QDA4E Subsystem User's Guide

QDA4E Disk Drive Controller SI55, SI56, and SI57 Disk Drives 5.25-Inch Disk Drive Quick Release DA123 Cabinet PB5500-9001 REV A



WARNING

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions contained in the Users Guide, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device when installed in accordance with the manufacturer's instructions, pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference, in which case the user, at his own expense, will be required to take whatever measures may be necessary to correct the interference.

NOTE

The following are trademarks of Digital Equipment Corporation:

BA23, BA123, DEC, DECnet, MASSBUS, MicroVAX I, MicroVAX II, MicroVMS, Q-bus, QBUS, SBI, UNIBUS, ULTRIX, VAX, VMS, and VT.

REVISION RECORD

Revision Number	Date	Description	EO Number
PB5500-9001	11/4/87	Initial Release	4630




TABLE OF CONTENTS

INTROL	DUCTION			
1.6 1.7 1.8 1.9	Manual Audience	•		1-2 1-2 1-3 1-5 1-5
1.10	configuration fittable actions	•	•	
PREINS	STALLATION			
2.3 2.4 2.5 2.6	AC Neutral	•	•	2-1 2-1 2-1 2-1 2-1
DA123	CABINET	•	•	3-1
SI55,	SI56, AND SI57 DISK DRIVES			
4.5 4.6 4.7 4.8 4.9 4.10 4.11 4.12 4.13 4.14	Performance Specifications Functional Specifications Reliability Specifications	•	•	4-4 4-5 4-5
	1.1 1.2 1.3 1.4 1.5 1.6 1.7 1.8 1.9 1.10 PREINS 2.1 2.2 2.3 2.4 2.5 2.7 2.8 DA123 8155, 4.1 4.2 4.3 4.4 4.5 4.1 4.1 4.1 4.1 4.1 4.1 4.1 4.1 4.1 4.1	1.1 Manual Overview 1.2 How to Read this Manual. 1.3 Manual Audience. 1.4 Materials Needed 1.5 System Introduction 1.6 Supported Host Computers,	1.1 Manual Overview 1.2 How to Read this Manual. 1.3 Manual Audience. 1.4 Materials Needed 1.5 System Introduction 1.6 Supported Host Computers,	1.1 Manual Overview 1.2 How to Read this Manual 1.3 Manual Audience. 1.4 Materials Needed 1.5 System Introduction 1.6 Supported Host Computers,

	4.17 4.18 4.19 4.20	Jumpers and Switches	-17
SECTION 5	5 1/4	INCH QUICK RELEASE	
	5.1 5.2 5.3 5.4	Quick Release Installation 5 Quick Release Cabling 5 Quick Release Operation 5 Write Protecting Drives 5	-1 -4 -4 -5
SECTION 6	QDA4E	CONTROLLER	
	6.1 6.2 6.3	Overview 6 Controller Specifications 6 Installation 6	-1 -5 -6
		6.3.1 QDA4E Switch Settings 6 6.3.2 RS232 Maintenance Terminal	
		Adaptor 6 6.3.3 Connecting to Drives 6	-10 -11
	6.4	WOMBAT Utilities 6	-13
		6.4.1 Starting Up WOMBAT 6 6.4.2 MicroVAX II 6 6.4.3 WOMBAT Menu Options 6 6.4.4 Master Menu Options 6 6.4.5 Disk Structure Menu Options 6 6.4.6 Disk Test Menu Options 6 6.4.7 Bad Block Management Menu Options 6 6.4.8 Shadow Options Menu Option 6 6.4.9 WOMBAT Disk Structure 6 6.4.10 Drives with Removable Media . 6 6.4.11 Error Recovery Procedures 6 6.4.12 Drive Shadowing 6 6.4.13 Computing Sectors Per Track . 6 6.4.14 WOMBAT Error Messages 6 6.4.15 WOMBAT Self-Diagnostics Initialization Procedures 6	-16 -16 -19 -19 -25 -36 -37 -36 -36 -37 -37
	6.5	MSCP Programming 6	-39
		6.5.1 Overview of MSCP 6 6.5.2 Controller Communications 6 6.5.3 Message Transmission 6 6.5.4 Data Transmission 6 6.5.5 Initialization 6 6.5.6 Registers 6 6.5.7 MSCP Commands	-40 -43 -45 -45

			6.5.9	Fatal	Contr	coller	Err	or	•	•	•	6-48
		6.6	QDA4E	Operati	ng Sys	stems		•		•	•	6-50
			6.6.1 6.6.2 6.6.3	Micro	VMS Or	erati	ng S	yst	ems	•	•	6-52
		6.7	Cache (Operatio	on .			•		•	•	6-60
			6.7.1 6.7.2 6.7.3 6.7.4 6.7.5 6.7.6 6.7.7	Read Cache Cache Cache Cache Cache Cache	Look-a Alloo Usage Assig Opera Disab	thead cation of the control of the c	Alg	ori	thm	•	•	6-60 6-61 6-61 6-62 6-62
		6.8	ESDI I	nterface	e			•		•	•	6-63
			6.8.1 6.8.2									
		6.9	Qbus I	nterface	. .			•		•	•	6-66
			6.9.1 6.9.2		rupts t Memo	ry Ac	 cess	•	• •	•	•	6-67 6-68
SECTIO	ON 7	MAINTI	ENANCE									
			Field 1	stics an Replacea l and Re	able U	Inits		•		•	•	7-1
				FIGUI	RES							
1-2. 1-3. 1-4. 1-5. 1-6. 1-7. 1-8. 1-9. 1-10. 1-11.	QDA4E QDA4E Two QI QDA4E Two QI QDA4E Two QI QDA4E Two QI QDA4E	BA123 Lowboy DA4E Lo Highbo DA4E Highbo DA4E Lo Highbo DA4E Highbo DA4E Highbo DA4E Highbo	Configury Cabine Cabine Cabine Cabine Cabine Cabine Cabine Capacita Capacit	n	igurat Config figura Confi Config roVAX AX Con CroVAX	ion . gurati tion. gurati gurati Confi figur Conf	on . ion. gura atio igur figu	tion		•	•	1-8 1-9 1-10 1-11 1-12

3-1.	Front View of DA123 with Four Drives	•	•	•				•	3-2
3-2.	DA123 Cabinet with Component Detail.	•	•	•	•	•	•	•	3-3
4-1.	Inside View of Drive		•	•	•	•	•		4-1
4-2.	Head Positioning Mechanism	•	•	•	•	•	•	•	4-8
4-3.	SI55 Switch and Jumper Locations		•	•	•	•	•	•	4-11
4-4.	SI56 and SI57 Jumper and Switch Loca	tio	ons	3.				•	4-12
4-5.	Interface Connector Location			•				•	4-13
4-6.	Bottom View of Drive Tray								
4-7.	Front View of Installed Drive		•						4-19
4-8.	Multiple Drive System Cabling			•		•			4-20
4-9.	Drive Cabling								
4-10.	Transition Panel Cabling						•	•	4-21
4-11.	Transition Panel Connector Detail								4-22
5-1.	Drive Cabling								
5-2.	Drive Slide Tray								
5-3.	Front of Quick Release Package								
5-4.	Inserting Drive								
5-5.	Installed Drive								
5-6.	Rear View of Quick Release Package .								
6-1.	QDA4E Switch Locations								6-7
6-2.	Cabled Drive								
6-3.	Fully Cabled QDA4E with Four Drives.						•	•	6-13
	~	•	-		-	-	-		. — -

SECTION 1 INTRODUCTION

1.1 Manual Overview

Section 7

This manual contains information to install, operate, and maintain the System Industries (SI) QDA4E subsystem. The information is presented in seven sections.

Section 1	Introduction: describes manual contents, standard features, capabilities, supported configurations, and options.
Section 2	Preinstallation Planning: describes site considerations including power and heat dissipation requirements.
Section 3	DA123 Cabinet: describes the installation and operation of the DA123 cabinet.
Section 4	SI55, SI56, and SI57 Disk Drives: describes the installation and operation of these disk drives.
Section 5	5 1/4 Inch Quick Release: describes the installation and operation of the 5 1/4 inch quick release mechanical package.
Section 6	QDA4E Controller: describes the installation and operation of the QDA4E disk drive controller.

The advantage of this manual is that it provides a single reference for most of the situations the user encounters. For more detailed information refer to the following System Industries and Maxtor publications:

Maintenance: provides maintenance, diagnostics, and troubleshooting information.

Part Number Title

PB5000-9200 ULTRIX Support Manual--SI Manual

1.2 How to Read this Manual

System Industries recommends that you read the entire manual to ensure as full an understanding of the QDA4E subsystem as possible.

1.3 Manual Audience

This manual is for use by System Industries selfmaintenance customers and field service technicians. System Industries assumes you have proper technical knowledge and equipment for the installation and maintenance of controllers, disk drives, cabinets, operating systems, and related software and equipment. If you are uncertain of this, contact a System Industries Customer Service Representative.

1.4 Materials Needed

No special tools or measuring instruments are needed to install and maintain the QDA4E, SI55, SI56, SI57, 5 1/4 inch quick release, and DA123 cabinet. Standard tools include flat head and phillips head screwdrivers, wire strippers and Allen wrenches,

1.5 System Introduction

The QDA4E subsystem is a disk and tape storage subsystem consisting of the following components:

- o QDA4E (Qbus Disk Adaptor for up to four ESDI drives)
- o SI55, SI56, or SI57 disk drive
- o DA123 Cabinet or space in host MicroVAX II

- o Cabling
- o Optional 5 1/4 inch disk drive quick release

See Figure 1-1 for a block diagram of the QDA4E subsystem.

1.6 Supported Hardware, Software, and Configurations

The QDA4E installs in DEC MicroVAX II computers running MicroVMS 4.1 and higher or ULTRIX version 1.2. QDA4E supports WOMBAT diagnostic software. No additional software is necessary.

Refer to Figures 1-2 to 1-12 at the end of this section for diagrams of the various configurations.




Figure 1-1. QDA4E Block Diagram

INTRODUCTION

The QDA4E supports the following drive and cabinet configurations:

	Max. Num. of Drives	Max. Num. of Quick Releases
Highboy Disk	4	2
Highboy Tape	4	2
Lowboy	4	2
BA123 or DA123	4	None

For new installations, the QDA4E supports the following SI and DEC cabinets:

6000-7288	Highboy 60 Hz disk cabinet
6000-7289	Highboy 50 Hz disk cabinet
6000-7285	Highboy 60 Hz tape cabinet
6000-7287	Highboy 50 Hz tape cabinet
4200-7333	Lowboy 60 Hz cabinet
4200-7334	Lowboy 50 Hz cabinet
	DA123 cabinet
H9647EA	DEC highboy
H9642JA	DEC lowboy

Filler panels mount on lowboys and highboys; select the panels from the following table:

Lowboy Cabinet

			Number of Quick Releases			With BA23/BA123 Rackmounted			
			1	2	:	L	2		
4200-1014 10.5 inch	Filler	Panel	2	2	-	L	1		
4200-1015 5.25 inch	Filler	Panel	1	0	2	2	1		

Highboy Tape Cabinet

	Number of Releases	Quick	With BA23/BA123 Rackmounted			
	1	2	1	2		
6000-1241 1.75 inch Filler Panel	1	1	-	_		
6000-1244 5.25 inch Filler Panel	1	-	1	_		
6000-1245 7 inch Filler Panel	2	2	-	_		
6000-1247 5.25 inch Filler Panel	_	_	1	1		

Each tape or disk highboy cabinet requires one 6001-7025 cable support bracket kit.

1.7 Cabling Configuratio1n

The QDA4E and ESDI drives can be cabled only in a daisy configuration.

The QDA4E supports single channel cabling to the drive.

ESDI cable restrictions state a maximum cable length of 3 meters or 9.8 feet. As such, the QDA4E supports only one quick release package with two SI55, SI56, or SI57 drives.

The following three sections describe specific cabling configurations.

1.8 Internal Cables for MicroVAX II BA123 and SI DA123 Installations

If the SI55, SI56, or SI57 drives are to be installed in a MicroVAX II, BA123, or SI DA123, select the appropriate cables from the table below. For add ons within the same cabinet, subtract the quantity of cables provided with the existing configuration from the quantity of cables that would be required to satisfy the addon configuration and enter the quantities on the lines provided.

9901-7982	Control Cabl	e, 1	Drive,	34	Conductor
9901-7983	Control Cabl	e, 2	Drive,	34	Conductor
9901-7984	Control Cabl	e, 3	Drive,	34	Conductor
9901-7931	Control Cabl	e, 4	Drive,	34	Conductor
9901-7930	Data Cable,	20 C	onducto	<u>-</u>	

Number of Disk Drives

	1	2	3	4
9901-7982 9901-7983 9901-7984	1	1	1	
9901-7931 9901-7930	1	2	3	1 4

1.9 MicroVAX II Cables and Transition Panels

If using the DA123 (external drives), for each QDA4E controller you need the 8650-7001 MicroVAX transition panel hardware kit.

Internal cables for connecting the QDA4E to the transition panel (DA123 for external drives) follow:

6005-7019 Cable Assy, 16 inch A Cable, ESDI CPU 6005-7020 Cable Assy, 16 inch B Cable, ESDI CPU

Select cables from the following table:

	1 Q1	DA4E	2	QDA4E
Number of Drives	1	2	3	4
6005-7019 6005-7020	1 1	1 2	2	2 4

For SI cabinets, no transition panel is required; the shielded external cables from the MicroVAX connect directly into the quick release.

External cables for BA23/BA123 to 5 1/4 inch quick release.

9901-7914 4 foot shielded cable

1.10 Configuration Illustrations

The following pages illustrate the various configurations.

1 CONTROLLER/4 DRIVES

CABLE PART NUMBERS:	DRIVE TYPE:
9901-7930 X 4 INTERNAL B	\$155 X4
9901-7931 X 1 INTERNAL A	\$156 X4
	S157 X4

BA123





Figure 1-2. QDA4E BA123 Configuration

1 CONTROLLER/2 DRIVES

CABLE PART NUMBERS: 6005-7019 X 1 A 16" ESDI CPU 6005-7020 X 2 B 16" ESDI CPU 9901-7914 X 2 4' EXTERNAL SHIELDED

DRIVE TYPE:

SI55 X2/ QUICK RELEASE PACKAGE X1 SI56 X2/ QUICK RELEASE PACKAGE X1 SI57 X2 QUICK RELEASE PACKAGE X1




NOTE:

AREA ABOVE 5 1/4 " QUICK-RELEASE ESDI PACKAGE MUST REMAIN EMPTY TO MAINTAIN FCC COMPLIANCE.

Figure 1-3. QDA4E Lowboy Cabinet Configuration

2 CONTROLLERS/4 DRIUES

CABLE PART NUMBERS: DRIVE TYPE: 6005-7019 % 2 A 16" ESDI CPU SI55 X4/ QUICK RELEASE PACKAGE X2 6005-7020 X 4 B 16" ESDI CPU SI56 X4/ QUICK RELEASE PACKAGE X2 9901-7914 X 4 4' EXTERNAL SHIELDED SI57 X4/ QUICK RELEASE PACKAGE X2



NOTE:

AREA ABOVE 5 1/4" QUICK-RELEASE ESDI PACKAGE MUST REMAIN EMPTY TO MAINTAIN FCC COMPLIANCE.

Figure 1-4. Two QDA4E Lowboy Cabinet Configuration

1 CONTROLLER/2 DRIVES


CABLE PART NUMBERS:

6005-7019 X 1 A 16" ESDI CPU 6005-7020 X 2 B 16" ESDI CPU

9981-7914 X 2 4' EXTERNAL SHIELDED

DRIVE TYPE:

SI55 X2/ QUICK RELEASE PACKAGE X1 SI56 X2/ QUICK RELEASE PACKAGE X1 SI57 X2/ QUICK RELEASE PACKAGE X1



NOTE

AREA ABOVE 5 1/4" QUICK-RELEASE ESDI PACKAGE MUST REMAIN EMPTY TO MAINTAIN FCC COMPLIANCE.

Figure 1-5. QDA4E Highboy Cabinet Configuration

2 CONTROLLERS/4 DRIVES

CABLE PART NUMBERS:

6005-7019 X 2 A 16" ESDI CPU

6005-7020 X 4 B 16" ESDI CPU


9901-7914 X 4 4' EXTERNAL SHIELDED

DRIVE TYPE:

SI55 X4/ QUICK RELEASE PACKAGE X2

SI56 X4/ QUICK RELEASE PACKAGE X2

SI57 X4/ QUICK RELEASE PACKAGE X2



AREA ABOVE 5 1/4" QUICK-RELEASE ESDI PACKAGE TO REMAIN EMPTY TO MAINTAIN FCC COMPLIANCE.

Figure 1-6. Two QDA4E Highboy Cabinet Configuration

1 CONTROLLER /2 DRIVES

CABLE PART NUMBERS: 6005-7019 X 1 A 16" ESDI CPU 6005-7020 X 2 B 16" ESDI CPU 9901-7914 X 2 4' EXTERNAL SHIELDED

DRIVE TYPE:

SI55 X2/ QUICK RELEASE PACKAGE X1 SI56 X2/ QUICK RELEASE PACKAGE X1 SI57 X2/ QUICK RELEASE PACKAGE X1




Figure 1-7. QDA4E Lowboy with MicroVAX Configuration

2 CONTROLLERS/4 DRIVES

CABLE PART NUMBERS:
6005-7019 X 2 A 16" ESDI CPU
6005-7020 X 4 B 16" ESDI CPU
9901-7914 X 4 4' EXTERNAL SHIELDED

DRIVE TYPE:

SI55 X4/ QUICK RELEASE PACKAGE X2 SI56 X4/ QUICK RELEASE PACKAGE X2 SI57 X4/ QUICK RELEASE PACKAGE X2




Figure 1-8. Two QDA4E Lowboy with MicroVAX Configuration

1 CONTROLLER /2 DRIVES

CABLE PART NUMBERS:

6005-7019 X 1 A 16" ESDI CPU

6005-7020 X 2 B 16" ESDI CPU

9901-7914 X 2 4' EXTERNAL SHIELDED

DRIVE TYPE: SI55 X2/ QUICK RELEASE PACKAGE X1

SI56 X2/ QUICK RELEASE PACKAGE X1 SI57 X2/ QUICK RELEASE PACKAGE X1




Figure 1-9. QDA4E Highboy with MicroVAX Configuration

2 CONTROLLERS/4 DRIVES

 CABLE PART NUMBERS:
 DRIVE TYPE:

 6005-7019 X 2 A 16" ESDI CPU
 \$155 X4/ QUICK RELEASE PACKAGE X2

 6005-7020 X 4 B 16" ESDI CPU
 \$156 X4/ QUICK RELEASE PACKAGE X2

 9901-7914 X 4 4' EXTERNAL SHIELDED
 \$157 X4/ QUICK RELEASE PACKAGE X2




Figure 1-10. Two QDA4E Highboy with MicroVAX Configuration

1 CONTROLLER /2 DRIVES

CABLE PART NUMBERS: 6005-7019 X 1 A 16" ESDI CPU 6005-7020 X 2 B 16" ESDI CPU

9901-7914 X 2 4' EXTERNAL SHIELDED

DRIVE TYPE:

SI55 X2/ QUICK RELEASE PACKAGE X1 SI56 X2/ QUICK RELEASE PACKAGE X1 SI57 X2/ QUICK RELEASE PACKAGE X1




Figure 1-11. QDA4E Highboy Tape with MicroVAX Configuration

2 CONTROLLERS/4 DRIVES

CABLE PART NUMBERS:

6005-7019 X 2 A 16" ESDI CPU 6005-7020 X 4 B 16" ESDI CPU

9901-7914 X 4 4' EXTERNAL SHIELDED

DRIVE TYPE:

\$155 X4/ QUICK RELEASE PACKAGE X2 \$156 X4/ QUICK RELEASE PACKAGE X2

S157 X4/ QUICK RELEASE PACKAGE X2




Figure 1-12. Two QDA4E Highboy Tape with MicroVAX Configuration

SECTION 2 PREINSTALLATION

Although System Industries sales representatives perform a thorough site analysis prior to installation, this section provides a useful summary and is helpful in planning future additions. Consider the following constraints when planning an installation of a QDA4E subsystem. The end of the section has a site installation checklist.

2.1 Power Requirements

Refer to the specific component section for power requirements. The QDA4E subsystem requires low power so that little presite consideration is necessary except for large installations or installations that have poor quality power or poor quality grounds.

2.2 Power Distribution Unit

For cabinet installations that may require an additional power distribution unit, the power distribution unit must be the new version, manufactured after May 1987, which is packaged in a slim line chassis that is mounted on the side of the cabinet.

2.3 AC Neutral

AC neutral must not be confused with protective or frame ground. The frame ground is used to prevent the build up of dangerous voltages on equipment and as protection for personnel—and to assure that any short circuit between a power phase and the cabinet draws enough current to trip the circuit's protective device immediately, rather than raising the potential of the equipment to a dangerous level. Additionally, this also prevents spurious noise from entering the line. Never connect AC neutral to the frame of any equipment or the protective ground (except at the building's main electrical service entrance).

Neutral and safety ground are often connected together by the NEMA receptacles or at the circuit breaker neutral bus bar. Attempts should be made to isolate neutral from safety ground in the circuit breaker box and ensure that conduit pipes are also isolated from other possible ground connections. Ideally, the equipment frame ground would be isolated from neutral and other ground sources all the way back to the building main grounding rod.

2.4 AC Earth Ground

An adequate earth ground must be maintained. If there is any question of inadequate grounding potential, perform an impedance test to ensure that it is less than 10 ohms. Refer to a good grounding reference for descriptions of ground impedance measurement techniques and grounding methods.

Typically, an earth ground suitable for computer equipment and peripherals consists of driving a 0.625 inch diameter copper rod into the earth to a depth of at least 12 feet. Since soil is quite variable in conductivity, chemicals such as salt or magnesium sulfate are added to the soil surrounding the rod to a depth of two feet. Periodic watering and chemical replenishment ensure an ongoing proper ground.

2.5 Air Flow

The SI55, SI56, and SI57 drives must have a linear airflow of 500 feet per minute across the PCB. Be sure that the front and rear vents of the cabinet are free from obstruction. Dress cables in cabinets to minimize airflow obstruction. Be cautious of enclosing MicroVAX II's or DA123's in confining environments that do not have adequate airflow.

2.6 Environmental Considerations

Refer to the component section for specific environmental considerations.

2.7 Other Site Considerations

For each storage subsystem, consider the following additional points:

- o Is there sufficient space for equipment and working area?
- o Are there any stairs, elevators, ramps, narrow doors or other obstacles that might impede delivery of equipment?
- o Are there proper fire and safety precautions?

- o Are the building requirements fulfilled: floor load, location of cables, AC outlets, and other equipment?
- o Are the peripherals and host computers sufficiently close so that the lengths of connecting cables do not exceed maximum limits?
- o Has the possibility of static been minimized? Antistatic flooring, chairs, wrist straps?
- o Is there an efficient workflow pattern to other work areas and an ease of visual observation of input and output devices?
- o Is there sufficient storage space for supplies necessary to the system?
- o Is there room for future expansion?

2.8 Site Preparation Checklist

The following page is a checklist to help with preinstallation:

Site Preparation Checklist

Task	Completion Date
Facility Construction	
Ease of Access to Equipment	
Floor Requirements	
System Calculation	
Current Requirement	
110 or 220 VAC	
Heat Dissipation	
Ground Impedance	
Voltage and Frequency Stability	
Power Cabling Routing	
Environmental Factors Met	
Temperature	-
Humidity	
Air Cleanliness	<u></u>
Vibration	
Accessible Telephone with 24 Hr Service	
Read Unpacking Caution Below	

UNPACKING CAUTION

Keep all shipping containers. Also, to avoid any question of liability, if any container arrives damaged or looks like it could be damaged, do not open it except in the presence of the shipping agent or representative—do not sign for any container that looks like it might be damaged. After opening a container inspect the equipment for damage and check contents against bill of materials—if anything is missing contact System Industries immediately.

SECTION 3 DA123 CABINET

The SI DA123 cabinet can be a valuable addition to a QDA4E subsystem. It provides an ideal environment for housing the SI55, SI56, and SI57 drives. The DA123 can also be used to house the BA23 MicroVAX I. A 350 watt power supply is incorporated in the DA123. A heavy duty fan provides airflow for the drives.

Figure 3-1 shows a front view of the DA123. Figure 3-2 provides a parts breakdown of a DA123 with four drives--the guide list follows:

	ITEM	DESCRIPTION
14 DISK Drive	2 3 4 5 6 7 & 8 9 & 10 11 12	DC Power Supply Cable Connector DC Power Supply Chassis AC Power Cable Tiedown AC Power Cable Input/Output Control Panel External Data/Control Cables Internal Data/Control Cables Fan




Figure 3-1. Front View of DA123 with Four Drives





Figure 3-2. DA123 Cabinet with Component Detail



SECTION 4 SI55, SI56, AND SI57 DISK DRIVES

The SI55, SI56, and SI57 disk drives are high capacity, high performance random access storage devices utilizing 8 non-removable 5 1/4-inch disks as storage media. Figure 4-1 provides cutaway view of drive with dimensions.




Figure 4-1. Inside View of Drive

These drives utilize the ESDI high performance 5 1/4-inch standard interface. Among the resultant benefits are a 15 Mbit/sec transfer rate, status and configuration reporting across the interface, and NRZ data transfer.

High performance is achieved through the use of a rotary voice coil actuator and a closed loop servo system utilizing a dedicated servo surface. The innovative rotary voice coil actuator provides average access times of 16 to 18 msec. and a track-to-track access time of 3 msec; performance usually achieved only with larger sized, higher powered linear The closed loop servo system and dedicated servo actuators. surface combine to allow state of the art recording densities in a 5 1/4 inch package. High capacity is achieved by a balanced combination of high areal new recording density, the incorporation of run-length-limited (RLL) data encoding techniques, and high density packaging techniques. Advanced electronic packaging techniques utilize miniature surface mount devices to allow all electronic circuitry to fit on one printed circuit board. Advanced 3380 Whitney type head flexures and thin-film sliders allow closer spacing of disks and therefore allow a higher number of disks in a 5 1/4-inch package. A unique integrated drive motor/spindle design allows a deeper HDA (head disk assembly) casting than conventional designs, thus permitting more disks to be used.

The electrical interface is compatible with the industry standards established by the ESDI committee. The size and mounting are also identical to the industry standard 5 1/4 inch minifloppy and Winchester disk drives, and they use the same DC voltages and connectors.

Key features of the SI55, SI56, and SI57 include the following:

- Same physical size and mounting as standard floppy disk drives.
- o Same DC voltages as standard floppy disk drives.
- o Rotary voice coil and closed loop servo system for fast, accurate head positioning.
- o Microprocessor controlled servo for fast access times, high reliability, and high density functional packaging.

- o ESDI compatible interface.
- o Thin film metallic media for higher bit density and resolution plus improved durability.
- o Single printed circuit board for improved reliability.
- o Automatic actuator lock.
- o Brushless DC spindle motor inside disk hub.
- o Microprocessor controlled spindle motor for precision speed control (+0, -.2%) under all load conditions.
- o Dynamic braking during power down cycle.
- o User selectable hard or soft sectors on SI56 and SI57.
- o Synchronization of spindle motors for parallel data transfer of multiple drives.
- o At power down, the heads automatically position over the non-data, dedicated landing zone on each disk surface. The automatic shipping lock solenoid also engages.

4.1 Environmental Limits and Physical Specifications

Ambient Temperature Operating: 50° to 113° F(10° to 45° C) Non-Operating: -40° to 140° F (-40° to 60° C)

Max Temperature Gradient Non-Operating or Operating: 18° F/hr.(10° C/hr), below condensation

Relative Humidity: 8 to 95% non-condensing

Maximum Elevation Operating: 10,000 ft. Non-Operating: -1000 ft. to 40,000 ft. Shock (inputs to frame of drive) Operating shock (all axes): 11 ms. pulsewidth (1/2 sine)...2.0GNon-operating shock (all axes): 11 ms. pulsewidth $(1/2 \text{ sine}) \dots 20.0G$ Vibration (inputs to frame of drive) Operating vibration (all axes) 5-25 hz, 0.006 inches P-P 25-500 hz, 0.20G peak acceleration Non-operating vibration (all axes) 5-31 hz, 0.02 inches P-P 31-500 hz, 1.0G peak acceleration DC Power Requirements +12V+ 5%, 1.5A typical, 4.5A max. (at power on) $+5V \pm 5\%$, 1.7A typical, 1.9A maximum +5V Maximum Ripple=50mV P-P +12V Maximum Ripple=120mV P-P Mechanical Dimensions Heat Dissipation 30 watts typical, 35 watts maximum. 4.2 Performance Specifications SI55 SI56 SI57 Capacity, unformatted

382.03

25.47

384.53

25.64

20,940

765.0

51.26

31,410

Per drive (Mbytes)

Per surface (Mbytes)

Per track (bytes) (minimum) 20,808

QDA4E SUBSYSTEM USER'S GUIDE S155, S156, AND S157 DRIVES

Capacity, typical format (512 byte sectors)			
Per drive (Mbytes) Per surface (Mbytes) Per track (bytes) Per sector (bytes)		280.17 18.68 15,872 512	23,552
Sectors/track	34	31	46
Transfer rate, Mbits/sec	10.0	10.0	15.063
Sectoring	Soft	Hard/Soft	Hard/Soft
Access time, msec, (maximum)			
Average [*] Track-to-track [*] Maximum	27 4 52	18 3 35	18 3 35
* Includes settling			

4.3 Functional Specifications

**			
Rotational speed (rpm) **	3600	3600	3600
Average latency (ms)	8.33	8.33	8.33
Recording density (bpi)	20,975	21,064	31,596
Flux density (fci)	13,980	14,043	21,064
Track density	1070	1070	1376
Cylinders	1224	1224	1632
Tracks	18,360	18,360	24,480
Sectors (512 bytes)	624,240	660,960	1,175,040
Data heads	15	15	15
Servo heads	1	1	1
Disks	8	8	8

^{**} Accurate to +0%, -0.2%

4.4 Reliability Specifications

MTBF	30,000 POH, typical usage
PM	Not Required
MTTR	30 minutes
Component Design Life	5 years

4.5 Error Rates

Soft read errors	10 per 10 ¹¹ bits read
Hard read errors*	10 per 10 ¹³ bits read
Seek errors	10 per 10 ⁷ seeks

^{*} Not recoverable within 16 retries

4.6 General Theory of Operation

The SI55, SI56, and SI57 disk drives consist of read/write and control electronics, read/write heads, servo head, head positioning actuator, media, and air filtration system. The components perform the following functions:

- o Interpret and generate control signals.
- o Position the heads over the desired track.
- o Read and write data.

4.7 Read/Write and Control Electronics

Drive Electronics are packaged on a single printed circuit board. This board, which includes two microprocessors, performs the following functions:

- o Reading/writing of data
- o Index detection
- o Head positioning
- o Head selection
- o Drive selection
- o Fault detection
- o Voice coil actuator drive circuitry.
- o Track 0 detection
- o Recalibration to track 0 on power-up
- o Track position counter

- o Power and speed control for spindle drive motor.
- o Braking for the spindle drive motor.
- o Drive up-to-speed indication circuit
- o Monitoring for write fault conditions
- o Control of all internal timing
- o Generation of seek complete signals
- o Run-Length-Limited (RLL) encoding/decoding
- o Data Separation
- o Address Mark Detection (Soft Sector)
- o Sector Detection (Hard Sector)
- o Spindle Synchronization

4.8 Drive Mechanism

A brushless DC drive motor contained within the spindle hub rotates the spindle at 3600 rpm. The spindle is direct driven with no belt or pulleys being used. The motor and spindle are dynamically balanced to insure a low vibration level. Dynamic braking is used to quickly stop the spindle motor when power is removed. The head/disk assembly is shock mounted to minimize transmission of vibration through the chassis or frame.

4.9 Air Filtration System

The disks and read/write heads are assembled in an ultra clean-air environment and then sealed within the module. The module contains an internal absolute filter mounted inside the casting to provide constant internal air filtration. A second filter, located on the bottom of the base casting, permits pressure equalization between internal air and ambient air.

4.10 Positioning Mechanism

The read/write heads are mounted on a head arm assembly which is then mounted to a ball-bearing supported shaft--see Figure 4-2. The voice coil, an integral part of the head/arm assembly, lies inside the magnet housing when installed in the drive. Current from the power amplifier, controlled by the servo system, causes a magnetic field in the voice coil which

either aids or opposes the field around the permanent magnets. This reaction causes the voice coil to move within the magnetic field. Since the head-arm assemblies are mounted to the voice coil, the voice coil movement is translated through the pivot point directly to the heads and achieves positioning over the desired cylinder.




Figure 4-2. Head Positioning Mechanism

Actuator movement is controlled by the servo feedback signal from the servo head. The servo head is located on the lower surface of the fourth disk from the top, where servo information is pre-written at the factory. This servo information is used as a control signal for the actuator to provide track-crossing signals during a seek operation, track-following signals during On cylinder operation, and timing information such as index, sector pulses (hard sector) and servo clock.

4.11 Read/Write Heads and Disks

The SI55, SI56, and SI57 employ thin-film head sliders and Whitney-type flexures. This configuration of sliders and flexures provides improved aerodynamic stability, superior head/disk compliance and a higher frequency response than conventional nickel--zinc heads.

The media utilizes thin metallic film deposited on 130mm diameter aluminum substrates. The coating formulation together with the low load-force/low mass Whitney type heads permit highly reliable contact start/stop operation. The metallic recording film yields high amplitude signals, and very high resolution performance compared to conventional oxide coated media. The metallic media also provides a highly abrasion-resistant surface, decreasing the potential for damage caused by shipping shock and vibration.

Data on each of the data surfaces is read by one read/write head. There is one surface dedicated to servo information in each drive.

4.12 Power-Up Sequence

DC power (+5V and +12V) may be supplied to the drive in any order, but +12 VDC is required to start the spindle motor. The motor power-up is controlled by the status of jumper JP6 on the drive electronics printed circuit board assembly--See Figures 4-3 and 4-4.

If jumper JP6 is open, the spindle power-up sequencing is initiated by the issuance of the control command START MOTOR.

When the spindle reaches full speed, the actuator lock automatically disengages and the heads then recalibrate to track 0. NOTE: Audible noise during the recalibration sequence is normal. Upon a successful recalibrate, READY and COMMAND COMPLETE status signals will be true. The unit will not perform any Read/Write or Seek functions until READY is true. (If after starting, 1000 rpm is not reached in 7 seconds, an automatic shutdown procedure is initiated; power to the spindle motor is shut off and the drive will not come READY.)

If jumper JP6 is shorted, the spindle power-up sequencing is initiated by the application of DC power. (When shipped, the JP6 jumper is shorted.)

4.13 Drive Selection

Drive selection occurs when the controller places the address of the drive to be selected on the three drive select lines. Only the selected drive will respond to the input signals, and only that drive's output signals are then gated to the The details of setting the drive selection jumper controller. are covered in the section Jumpers and Switches.

DRIVE SELECTED	DRIVE SELECT 3	DRIVE SELECT 2	DRIVE SELECT 1
NONE 1 2 3 4 5 6 7	0 0 0 0 1 1 1	0 0 1 1 0 0	0 1 0 1 0 1

4.14 Drive Termination

If more than one drive is used in a system, the termination resistor packs (RN 13 and RN 14 for SI56 and SI57) have to be removed in all but the last drive in the string. Figures 4-3 and 4-4 show the location of RN 13 and RN 14.

4.15 Electrical Interface

The interface to the SI55, SI56, and SI57 drives can be divided into four separate categories, each of which is physically separated.

- o Control signals
- Data signals
- DC power
- Auxiliary Signals




Figure 4-3. SI55 Switch and Jumper Locations




Figure 4-4. SI56 and SI57 Jumper and Switch Locations

All Control lines are digital in nature (open collector TTL) and either provide signals to the drive (input) or signals to the host (output) through interface connection J1/P1. The data transfer signals are differential in nature and provide data either to (write) or from (read) the drive through J2/P2.

4.16 Physical Interface

The electrical interface between the SI55, Si56, and SI57 and the host controller is through four required connectors and one optional connector. Figure 4-5 provides the locations of these connector/ports.

J1 - Control signals (multiplexed)
 J2 - Read/write signals
 J3 - DC power input
 J4 - Frame ground
 J6 - Optional spindle synchronization connector




Figure 4-5. Interface Connector Location

4.17 Jumpers and Switches

The following jumpers and switches are normally preset at the factory; many of them are not used or should not be changed. Factory settings for all jumpers are provided in case a jumper is accidently removed. See Figures 4-3 and 4-4 for jumper/switch locations.

Drive Address Selection Jumper

In multiple drive configurations, it is necessary to configure each drive with a unique address. A maximum of seven drives are permitted per single host controller—a maximum of four for the QDA4E. The address for the drive is determined by the location of the drive select jumper plug (DS 1) shown in Figure 4-3 and 4-4. The pins are numbered from 1 to 7 from left to right. The upper row of pins are common. As shipped from the factory, the drive is configured as logical unit 1. Removing the jumper entirely is equivalent to a "no select".

Drive Select Number	Jumper Installed
DV*0	DS1
DV*1	DS2
DV*2	DS3
DA*3	DS4
DV*4	DS5
DV*5	DS6
DV*6	DS7

Data Head Selection Jumpers

These jumpers determine the number of data heads enabled for the particular model of drive. They are set at the factory as follows:

SI55

	_	IP6 IN	JP7 IN	JP8 IN		
SI56 a	nd SI	57				
	J	TP32 IN	JP33 IN	JP34 IN	JP35 IN	JP36 OUT

Write Protect Selection Jumper

This jumper (JP9 for S55 and JP14 for SI56 and SI57) is the write project jumper. When the jumper is installed, the drive is write-protected and can only be read. As shipped from the factory, the jumper is removed.

Sequential Spindle Motor Spin-Up Jumper

The spindle motor spin-up jumper (JP10 for SI55 and JP6 for SI56 and SI57) allows a string of drives to be started sequentially by the controller. When the jumper is installed, the drive automatically spins up as soon as power is applied. If it is removed, the drive is started by issuing the appropriate command from the controller. As shipped from the factory, jumper is installed.

Off-Track Error Reporting Jumper--SI55 Only (JP11)

The off-track error jumper (when installed), allows the drive to generate ATTENTION with write fault status when an off track condition occurs during a write operation. As shipped from the factory, jumper JP11 is installed.

Test Jumpers

These jumpers provide access to certain test signals. The specific signals and the normal factory settings are shown below:

For SI55

JUMPER	FACTORY SETTING	NOTES ON FUNCTION
JP12	IN	USED FOR PHASE MARGIN TESTING; ECL LEVEL DATA CLOCK OUTPUT= PIN 1; INPUT= PIN 2.
JP13	IN	USED FOR PHASE MARGIN TESTING; ECL LEVEL DATA OUTPUT=PIN 5; DATA INPUT=PIN 4.
JP14	OUT	NEVER SHOULD BE JUMPERED; TWO TEST POINTS: RAW DATA FROM PULSE DETECTOR = PIN 6 GROUND = PIN 7
JP15	OUT	NOT USED
JP16	IN	ENCODED WRITE DATA; TTL OUTPUT OUTPUT = PIN 10; INPUT =PIN 11

NOTE: Pins E1 and E2 are ground pins.

For SI56 and SI57

JUMPER	FACTORY SETTING	NOTES ON FUNCTION
JP1	IN	Encoded write data;TTL
JP2	IN	Used for phase margin testing; ECL level clock. Output=PIN 18. Input=PIN 19.
JP3	IN	Used for phase margin testing; ECL level data. Output=PIN 20. Input=PIN 21

Sector Configuration Jumpers--SI56 and SI57

Jumper JP31 selects the mode of operation. Jumper JP31 installed configures the drive as a soft sector drive; removed configures the drive as a hard sectored drive.

Jumpers JP16-29 allow the user to configure the drive's hard sector size. The sector size can range from a minimum of 123 to a maximum of 15,625 (SI57) and 10,470 (SI56) unformatted bytes per sector, with 1 byte granularity.

The hard sector configuration jumpers are encoded in a binary fashion with JP16 being the LSB (least significant bit) and JP29 being the MSB (most significant bit). An installed jumper equates to a one.

SI56 and SI57 are set at the factory as follows:

	SETTINGS		
JUMPER	56	<u>57</u>	
JP16	IN	OUT	
JP17	IN	IN	
JP18	OUT	TUO	
JP19	OUT	IN	
JP20	OUT	OUT	
JP21	IN	IN	
JP22	OUT	OUT	
JP23	IN	IN	
JP24	OUT	OUT	
JP25	IN	IN	
JP26-JP31	OUT	OUT	

Jumper JP30, if installed, disables setting the hard sector size over the ESDI interface.

4.18 Installing the SI55, SI56, and SI57

Normally the drives are installed before shipment. However, for add-ons or selfmaintenance customers, the following instructions apply.

If installing the drive in BA123 or DA123 cabinet, attach a drive plate (part number 9901-7922) to the drive base--this is also usually done at factory. Four screws attach the plate to the drive--use a #1 flathead screwdriver--see Figure 4-6. The drive then slides into the cabinet; a locking tang on the cabinet chassis between the slides locks the drive in place. Pushing down on the tang allows you to remove drive--see Figure 4-7 for front view of drive and locking tang.

If installing drives in quick release package, refer to the 5 1/4 Inch Quick Release section for instructions.




Figure 4-6. Bottom View of Drive Tray




Figure 4-7. Front View of Installed Drive

4.19 Cabling Drives

The drives cable directly to the transition panel or to the QDA4E--Figure 4-8 provides a cabling diagram. Every drive has the following cables:

- o A Cable--wide 60 conductor cable
- o B Cable--narrow 26 conductor cable
- o Power supply cable--white Molex connector
- o Ground wire

In DA123 cabinets the A and B cable attach to the transition panel. In MicroVAX II's they attach to the QDA4E. Power supply cables have white Molex connectors that hang off the power supply. The ground wire screws into grounding screws. Figure 4-9 provides a detail of cabling the drive. Figure 4-10 provides a transition panel detail. Figure 4-11 provides a connector detail.




Figure 4-8. Multiple Drive System Cabling




Figure 4-9. Drive Cabling




Figure 4-10. Transition Panel Cabling




Figure 4-11. Transition Panel Connector Detail

Since there are so many ways of cabling the various drive and controller configurations the best way of cabling is to use the illustrations and cable configurations in the Introduction as well as Figures 4-8 to 4-11. If using quick release refer to quick release section and if using DA123 refer to DA123 section for more information.

4.20 SI55, SI56, and SI57 Operation

No user operation of the drive is necessary. The red LED on the front of the drive indicates read/write activity and is an excellent indicator that drive is working/active. Never remove cables between drive and controller or power down any part of the subsystem when the indicator is lit.

SECTION 5 5 1/4 INCH QUICK RELEASE

The 5 1/4 quick release mechanical package is a convenient means of mounting SI55, SI56, and SI57 drives in higboy or lowboy cabinets. Each quick release package holds up to two drives.

5.1 Quick Release Drive Installation

Typically drives are installed in quick release packages at the factory. Quick release packages install in cabinets as The following instructions are helpful for rack mounts. add-ons and selfmaintenance customers.

The cables, listed below, attach to the drive as shown in Figure 5-1. Note that the two gray connectors are directional -- the connector end with the arrow ingrained in the plastic shell aligns over the notch in the drive port; this also allows the cable to lift up and away from plate. The power supply connector attaches to the port on the bottom left rear of drive, as you face front of drive. A ground wire attaches to recessed spade terminal.

The cables and connectors follow:

- o A Cable--wide 34 pin cable
- B Cable--narrow 20 pin cable
- Power supply cable--with white Molex connector
- o Ground wire

The drive mounts on the slide tray by four screws provided with tray--Figure 5-2.

The drive and slide tray insert into quick release package; open quick release cover by pressing and releasing top and center of cover--Figure 5-3. Align drive with slides and push in--Figure 5-4. The connectors mate automatically -- drive release bar clicks into place. 5-5 shows the drive installed.




Figure 5-1. Drive Cabling





Figure 5-2. Drive Slide Tray






Figure 5-4. Inserting Drive




Figure 5-5. Installed Drive

5.2 Quick Release Package Cabling

Follow the illustrations and cable configurations in the Introduction--Figure 5-6 provides a rear view of quick release package with transition panel, power supply port, and power switch.

5.3 Quick Release Operation

Operation of the quick release consists of drive removal and write protecting drives.

Power off the quick release package by pushing the 1 position of the power switch on rear of quick release package.

A tricolor LED guides you through drive removal; remove drive as follows (if you aren't replacing drive immediately, reconfigure system for one less drive):

o Press release bar; handle lifts slightly.




Figure 5-6. Rear View of Quick Release Package

- o Lift handle; LED changes color from red to amber--drive is stopping which takes about 15 seconds.
- o When LED turns green, pull drive from quick release.

5.4 Write Protecting Drives in Quick Release

Each drive can be write protected as follows:

o Open quick release cover--press and release top and center of front.

- o The WRITE/PROTECT switch on the right and left are for unit 0 and unit 1 drives. Press the switch on the right; the LED inside the switch goes on and the drive on the right is now protected from having data written to it.
- o Press the switch again and the LED goes off and the drive is no longer write protected.

SECTION 6 QDA4E CONTROLLER

6.1 Overview

The QDA4E is a high performance quad height interface to ESDI compatible disk drives. The QDA4E can support up to four ESDI drives and features a one megabyte cache memory, 24 megabit per second throughput, command queueing, overlapped seeks, and implements DEC's Mass Storage Control Protocol (MSCP).

The QDA4E flexibly couples disks of any size and data rate to the MicroVAX II operating system without software modification. Comprehensive on-board interactive formatting and diagnostic firmware provides engineering support for MicroVAX II.

The ESDI Interface

The ESDI or Enhanced Small Device Interface is a high performance interface suitable for smaller high performance hard disk drives.

One Megabyte Disk Cache

All data is read from and written into the cache. Data transfers from the cache are approximately 2.5 ms compared to 30-38 ms for typical drive access times—up to a 93 percent reduction in access time. The cache memory acts as a track buffer and can provide data at rates of 3 megabytes per second.

Four ESDI drives at 3 Megabytes per second

The QDA4E controller supports up to four ESDI interface drives. All four drives can have different storage capacities, speed and data rates. This allows the user to match drive characteristics with applications.

No Sector Interleaving

The QDA4E stages all data from the disk drive through the cache memory. This ensures that all data can be transferred at full disk speed over the disk interface and at maximum speed over the Qbus without incurring data late errors. Disk and Qbus transfers are performed simultaneously to minimize access times.

Read Look-ahead

The QDA4E allows the user to program the controller to perform 'look-ahead' reads in anticipation of data requests. Whole tracks or more can be read into the cache ensuring that data is ready for the host computer the instant it's needed.

Virtual Units

The QDA4E also allows the user to partition each drive into virtual units which are addressed by the host as individual drives. Each virtual unit can be any size up to the size of the entire drive with up to 16 virtual units assigned to each controller. Each virtual unit can be further partitioned under the host operating system.

Block mode DMA and DMA Throttle

With Block Mode DMA, the QDA4E interleaves address references with bursts of data--almost doubling Qbus throughput. The QDA4E fully conforms with Qbus Block Mode DMA protocol. With non block mode memory, the QDA4E automatically reverts to burst mode DMA.

After every 16 word DMA transfer there is a 4 microsecond delay to service any pending interrupt or DMA requests from other devices. If a DMA request occurs a 'DMA throttle' releases the Qbus after 8 words to prevent data loss from other DMA devices.

Drive Shadowing

The QDA4E offers the user the option of drive shadowing. Data integrity is further improved by writing the same data to two drives simultaneously. In certain circumstances greater data throughput may be achieved because the controller can decide which drive has its heads positioned closest to the required data, and schedule a read on that drive.

Seek optimization and Overlap

The QDA4E can queue up to 32 commands. The optimum order of execution is dynamically computed according to the strategy selected by the user. With multiple disk drives seeks are initiated simultaneously, further improving performance.

Error Checking and Correction

The QDA4E uses a 48-bit ECC polynomial with an 11-bit correction span for error detection and correction. The QDA4E tries up to 10 times to correct an error before reporting the fault to the host system.

Dynamic Bad Block Replacement

QDA4E dynamic bad block replacement and error correction always presents error free 'perfect media' to the host computer. During normal operation the controller dynamically replaces any blocks it detects as bad with an alternative block from a replacement block pool. Blocks with hard errors are replaced but the data in them flagged 'forced error'. This indicates to the host that though the data in these blocks is bad the blocks themselves are now good. All bad block replacement is completely invisible to the host computer.

Statistics Recording

The QDA4E records statistics such as the number of reads and writes, cache hits and misses, and other important information for each drive. The user can interrogate the drive for this information according to application specific performance requirements.

Write Protect

A connector is provided to which the user can connect one write protect switch per drive.

MSCP Emulation

The QDA4E communicates with the host through a simple register pair to memory resident 'command packets'. Disk geometry factors such as sectors, heads, cylinders, and disk capacity are invisible to the host computer. The QDA4E accepts 32 bit binary block numbers and converts them to physical disk addresses, allowing any size disk to be fully accessed by any program without software modification.

Qbus Interface

Originally introduced in 1975 by Digital Equipment Corporation to support the LSI-11 CPU range, the Qbus

architecture has evolved in speed and function to the point where it now outperforms most small computer bus systems. The QDA4E fully supports MicroVAX II Qbus CPU.

On-Board WOMBAT Utilities

This is an interactive formatting and diagnostic utility contained within the QDA4E firmware. An on-board serial connection allows WOMBAT to be run using an ASCII terminal. This permits disk formatting and maintenance operations to be carried out with minimal additional hardware.

WOMBAT can also load a simple console communication program into the host computer's memory or it can be invoked on system power-up. No external software, media, or program loading device is required in maintenance of the QDA4E or its attached disk drives. WOMBAT is always available independently of the host CPU type or the operating system environment.

WOMBAT Formatter

WOMBAT initializes a fresh disk drive by writing sector addresses and zero data blocks through the entire recording surface. WOMBAT prompts the user at the terminal to support parameters such as drive geometry (cylinders, heads and sectors) and various other options. This data is stored twice in special reserved areas of track zero and retrieved by a simple homeseek-read sequence at each power-up. No special PROMs or switch settings are required to fully characterize the connected disk drives.

WOMBAT Self Diagnostics

The QDA4E contains a comprehensive set of self diagnostic procedures which are executed automatically on power-up. Failure is indicated by a flashing red LED and a fatal error status which is deposited in the SA register.

WOMBAT Interactive Diagnostics

Terminal oriented engineering utilities contained within the WOMBAT firmware include a continuous read/write/seek exerciser, a disk surface pattern tester and a bad block replacement routine.

Sequential Spin-up

The QDA4E controller starts ESDI drives in a sequential manner, providing they have the "motor control" option enabled. This is done to minimize start-up current surge.

6.2 Controller Specifications

Bus Interface (Qbus): MicroVAX II

Qbus Loads: 1 DC, 1 AC

Transfer Mode: Block Mode DMA

Memory Address: 4 megabyte capacity (22-bits)

Software Emulation: DEC MSCP

Command Buffer: Up to 32 commands capacity

Disk Cache Size: 1 megabyte (with parity)

2.5 ms for cached reads

Transfer Rate: 3.0 megabyte per second maximum

non-interleaved

Base (CSR) Address: Switch selectable

CSR MicroVAX

17772150 20001468 17760334 200000DC 17760340 200000E0 17760344 200000E4 17760354 200000EC 17760360 200000F0

Vector Interrupt: Programmable

Interrupt Level: Switch selectable

4, 5, 6, or 7

On-board Bootstrap: Switch selectable

Disable or enable at 17773000 or

17771000

Power Requirements: +5 VDC @ 3.5A typical

Dimensions: Standard quad board

LED Indicators:

Red: Fatal Error

Green: Access in progress

TTL Outputs: Disk access in progress

RS-232 Output: Data transmitted to terminal 9600 baud

RS-232 Input: Data received from terminal 9600 baud

Drive Support: 1 to 4

Interface: ESDI (Enhanced Small Device

Interface). Hard and soft sectored

drives are supported.

Connectors: 1 34-pin control connector

4 20-pin data connectors

1 10-pin Write Protect front panel

connector

Cylinders: 4096 maximum

Heads: 16 (maximum)

Bytes/Sector: 512

Sectors/Track: 255 (maximum)

6.3 Installation

Installation consists of setting the various switches, installing card, and cabling to drives.

6.3.1 QDA4E Switch Settings

The QDA4E can be configured by selecting the appropriate switch combination. The switch positions for base (CSR) address, automatic bootstrap select, interrupt priority, and CPU type are given below. Note that there is no interrupt vector switch as this is set automatically by the operating system software. The locations for switches are shown in Figure 6-1.




Figure 6-1. QDA4E Switch Locations

Base CSR Address

The base CSR address should be selected according to the rules set out in the section on operating systems. It is important to locate the controller correctly to enable to the operating system to identify the device and its type. The following is a table of all the possible CSR addresses for MicroVAX systems. Note that W3 is used and is much the same as W2 except that it cannot be the 4th controller if there are no DZV11's--if there are, it can be the third.

-----ADDRESS-----*SW2 STATUS CSR MICROVAX II SW2-10 -ON 17772150 20001468 SW2-9 -ON 17760334 20000DC SW2-8 -ON 200000E0 17760340 -on SW2-7 17760344 200000E4 SW2-6 -ON 17760354 200000EC SW2-5 -ON 17760360 200000F0

BASE (CSR) ADDRESS - SWITCH (SW2)

*NOTE - Only one of the above six switches should be on at any one time.

QDA4E FACTORY SWITCH SETTINGS

Automatic Bootstrap

The on-board bootstrap is provided to allow systems to boot from the controller. The switch settings for the bootstrap are given below. The bootstrap must be disabled before being installed in a MicroVAX.

AUTOMATIC BOOTSTRAP SELECT - SWITCH (SW2)

*SW2 STATUS BOOTSTRAP FUNCTION				
SW2-3 - ON SW2-2 - ON SW2-1 - ON	On board bootstrap disabled (for MicroVAX) On board bootstrap enabled at address 17773000 On board bootstrap enabled at address 17771000			

*NOTE - Only one of the above three switches should be on at any one time--SW2-2 is normally on. Switch 4 is unused.

CPU Type

This feature is used in combination with the automatic bootstrap. If the bootstrap is enabled the appropriate CPU type must be selected. CPU type 3 is selected to run WOMBAT from the on-board terminal port when the system is initialized.

CPU TYPE - SWITCH (SW3)

CPU TYPE	*SW3-4	SW3-3
0 (MicroVAX)	ON	ON
3 (No CPU)	OFF	OFF

Interrupt Priority Setting

Interrupts suspend program execution while the processor starts the device service routine at a vector address input from the requesting device.

Interrupts are serviced according to device priority. Device priority can be determined in two ways. These are termed 'Position Defined' and 'Distributed' arbitration. Positioned Defined arbitration give priority to those devices which are electrically closest to the processor. Distributed arbitration implements priority according to the priority levels set on the device hardware. When devices with equal priority generate an interrupt, the processor gives

preference to the device which is electrically closest. A previous bus transaction must have been completed before another can be commenced.

INTERRUPT PRIORITY - SWITCH (SW1)

SW1-2	SW1-3	SW1-4	PRIORITY
ON	ON	ON	4
ON	ON	OFF	5
ON	OFF	ON	6
OFF	OFF	ON	7

NOTE - SW1-1 is unused. Switch is set at factors to 4.

6.3.2 RS232 Maintenance Terminal Adaptor

The optional RS232 maintenance terminal adaptor allows the controller to connect to an ASCII terminal. It consists of a 10-pin socket on a 10-conductor cable that terminates with either a male DB25 connector or a female DB25 connector.

The communication format is:

ASCII RS232 9600 Baud, 8 Data Bits, 1 Stop Bit, no parity.

Note that if a VT220 terminal is used, the communication format must be set up for space parity rather than no parity.

If normal disk access is attempted with this cable connected to a terminal, garbage appears on the terminal due to the shared RS232 Output/Access Light Function. This is normal.

J6 Pins DB25S Pins Function

7 7 RS232 Enable
8 2 RS232 Input
3 3 RS232 Output
4 7 Ground

RS232 MAINTENANCE TERMINAL ADAPTOR

6.3.3 Connecting to Drives

The QDA4E supports up to four ESDI drives each of which can be of different size, speed and data rate. The procedure for connecting multiple drives is simple. Data cables are connected in radial fashion to connectors on the QDA4E. There is no particular order of connection. The control cable is daisy chained between each drive with the last drive being terminated as described in Section 4.

Pre-configured data and control cable kits for one to four drives are supplied to minimize difficulties when installing the QDA4E or where adding additional drives.

Cables should be checked for correct installation by ensuring that Pin 1 on the connector aligns with Pin 1 on the cable. Pin 1 is generally indicated by the arrow mark on the connector shroud. If there is no shroud Pin 1 is the top right pin with the connectors facing you. Pin 1 on the cable is usually indicated by a colored stripe on the appropriate cable.

Figure 6-2 shows a drive (without required base plate for illustration) cabled to QDA4E. Insert drive, attach the power supply connector, and attach the ground wire to cabinet chassis ground screw. Figure 6-3 shows a QDA4E installation with four drives.

Correct drive termination is also necessary. Consult the disk drive section for the appropriate termination procedure.




Figure 6-2. Cabled Drive




Figure 6-3. Fully Cabled QDA4E with Four Drives

6.4 WOMBAT Utilities

6.4.1 Starting Up WOMBAT

WOMBAT provides a controller resident means of formatting, testing and maintaining the drive and controller subsystem. All WOMBAT functions are menu driven and are designed to simplify the process of structuring, formatting and testing drives.

WOMBAT can be invoked using any of the following methods:

- (a) By selecting the 'W' option during the system bootstrap operation if the QDA4E boot is enabled. This allows disk testing and diagnostics to be performed from the user console. The console link is formed by a communication program which WOMBAT downloads into main memory.
- (b) By entering 254, (MicroVAX II) from the console terminal using ODT. This allows disk testing and diagnostics to be performed from the user console. The console link is formed by a communication program which WOMBAT downloads into main memory.
- (c) By connecting a 9600 baud auxiliary terminal to the QDA4E Maintenance Connector and entering 260g from the console terminal using ODT. This allows disk testing and diagnostics to be performed from an auxiliary terminal which communicates directly with on-board WOMBAT utilities.
- (d) By connecting a 9600 baud auxiliary terminal to the QDA4E Maintenance Connector and configuring the switches on the PCB to automatically run WOMBAT. (CPU independent.)

Runs WOMBAT independently of, or without, a CPU for controller testing or engineering purposes if necessary. First disable the controller bootstrap and set the CPU type to "No CPU". Connect a 9600 baud terminal to the Front Panel connector and position the controller in the backplane. WOMBAT will be invoked automatically on power-up or by pressing reset. If no CPU is present then the backplane must be correctly terminated and a bus initializing signal (BINIT) must be generated.

To resume normal operation the configuration switches must be reset as required. Note that setting the CPU type to 3 will make the controller completely unavailable to the host CPU.

(e) By depositing 272, into the IP register a user-written communication routine can communicate directly with WOMBAT. This call does not outload any communication routine into the host memory.

The following table summarizes the procedures that invoke WOMBAT:

WOMBAT INITIALIZATION PROCEDURES

ACTI	ON	CONDITIONS	CONTROL
*OCTAL 000254 000260 000272	*HEX 00AC 00BO 00BA	ODT ODT ON-LINE	MicroVAX Console Aux Terminal on Maint. Connector User Communication Program
W		Via Boot	System Console
SW2-3 ON SW3-3,-4		CPU Independent	Aux Terminal on Maint. Connector

^{*} These codes are deposited into the IP register.

The procedures for invoking WOMBAT on the MicroVAX II are given below. WOMBAT can be stopped by simply re-booting the system.

6.4.2 MicroVAX II

The following details the procedure for invoking WOMBAT in a MicroVAX II using ODT.

Halt the CPU at the end of its start-up diagnostics by turning on the "halt enable" switch at the back of the CPU.

When it halts, do the following:

D/P/W	20001F40	20	<pre>(enable QBUS access to memory)</pre>
D/L	20088008	80000002	(set up the appro- priate QBUS map entry)
D/W	'CSR'	AC	<pre>(ask WOMBAT to load the communications program into memory)</pre>
S	400		(start the program)

Where the QDA4E addresses are as follows:

CPU CSR	MicroVAX II
17772150	20001468
17760334	20000DC
17760340	20000EC
17760344	200000E4
17760354	200000EC
17760360	200000F0
	17772150 17760334 17760340 17760344 17760354

6.4.3 WOMBAT Menu options

When WOMBAT is invoked it will display an announcement and then print a list of all drives and units and prompts for the drive number on which to perform operations.

QDA4E		WOMBA	AT Version	n: 2.x	where x is the firm- ware revision level								
UNIT	DRIVE	OFFSET	SIZE	WRITE	STATUS								
0	0	34	20000	LATE	AVAIL								
DRIVE	CYLS	HEADS	SECTORS	BLOCKS		FAIR	STATUS						
0	1224	5	34	291312		24	SPUN UP						

Drive number:

Enter the drive number (zero on a single-drive system). WOMBAT will then display the Master Menu options.

** Master Menu ** 1

- 1 Structure Disk
- 2 Test Disk
- 3 Manage bad blocks
- 4 Display error
- 5 Shadow options

Select an option by typing the option number followed by RETURN. Options 1 through 3 and 5 will provide sub menus while option 4 displays the last controller detected fatal error. To return to the master menu from a sub menu type RETURN.

To exit from the master menu to the announcement (to select a different drive) type RETURN. WOMBAT will not allow you to do this before verifying whether the disk structure data has been written to disk. 'NO' is the default value.

6.6.4 Master Menu Options

Option 1 - Structure Disk

Selecting this option causes a sub menu to be displayed as follows:

- ** Disk Structure Menu **
- 1 Create Disk Structure
- 2 Format Disk
- 3 Write Disk Structure
- 4 Update Header Blocks
- 5 Display Disk Structure
- 6 Change Unit number
- 7 Display Statistics

Option 2 - Test Disk

Selecting this option causes a sub menu to be displayed as follows:

- ** Disk Test Menu **
- (! means all data on Disk destroyed)
- 1 Read All Disk (preserves all data)
- 2 ! Write Disk !
- 3 ! Pattern Test!
- 4 ! Random Writes!
- 5 ECC Validation
- 6 Read Physical Block

- 7 Display Error Statistics
- 8 Zero Error Statistics
- 9 Test Cache, RAM & ROM

Option 3 - Manage Bad Blocks

Selecting this option causes a sub menu to be displayed as follows:

- ** Bad Block Management Menu **
- 1 Manually Replace Bad Block
- 2 Automatically Replace Bad Blocks
- 3 Display Replaced Bad Blocks
- 4 Enter defect map
- 5 Get defect map from drive

Option 4 - Display Error

Selecting this option causes WOMBAT to display a message which explains the most recent controller detected fatal error. The occurrence of an error is indicated by the controller hanging with the red LED flashing. If no error has occurred this option produces a meaningless error message.

Option 5 - Shadow Options

Selecting this option causes a sub menu to be displayed as follows:

- ** Shadow Options Menu **
- 1 Shadow copy option
- 2 Set shadow units
- 3 Reset shadow units
- 4 Copy unit to unit
- 5 Compare unit to unit

Setting Up a New Disk

The procedure for structuring a new disk is as follows:

- 1. Create disk structure.
- 2. Format the disk.
- 3. Write the disk structure.

- 4. Replace bad blocks using manufacturer's media defect map.
- 5. Pattern Test the disk.
- 6. Replace bad blocks detected by pattern test.

Once these 6 steps have been undertaken the host operating system may use the disk.

6.4.5 Disk Structure Menu Options

- ** Disk Structure Menu **
- 1 Create Disk Structure
- 2 Format Disk
- 3 Write Disk Structure
- 4 Update Header Blocks
- 5 Display Disk Structure
- 6 Change Unit Number
- 7 Display Statistics

Option 1 - Create Disk Structure

The Create Disk Structure option must be performed when a new disk is connected to the QDA4E. This allows the disk geometry, controller wide tuning parameters, and virtual unit structure to be specified. The virtual unit structure allows a single large drive to appear to the host operating system as multiple drives.

This option enters an interactive question and answer dialogue which specifies the disk structure. WOMBAT displays either the current or the default value for a parameter and gives you the option of accepting or changing it to a new value. To accept the displayed value, hit RETURN. To change it, type in the new value followed by RETURN. If WOMBAT detects improper values it will issue a warning.

The create disk structure dialogue is divided into three parts.

- 1. The drive structure specification, which describes the physical geometry of the drive.
- The unit structure specification, which is executed once for each virtual unit defined where the size of each

unit and unit specific parameters is described.

3. The controller wide tuning parameters, where read look ahead, and command queue size are specified.

Drive structure specifications

Cylinders:

The number of cylinders on the drive. There normally is no need to change this value. Note that this will be one less than the number specified by the manufacturer - this is to preserve the manufacturer recorded defect list on the last cylinder. You may override this option and the last cylinder, but the manufacturer's defect list will be destroyed during formatting.

Heads:

The number of heads on the drive is displayed. There is normally no need to change this value.

Full sectors/track:

Hard sectored drives report the correct value For soft sectored drives, the required value will have to be computed and entered. See Sectors Per Track.

Track spiralling:

Track spiralling improves disk performance on transfers over more than one track. Sector zero on track is offset by a nominated factor to allow select and positioning before sector zero on the track is reached. The recommended factor is four.

Optimization:

The seek optimization strategy can be either:

- O (None) No optimization done. First request executed. This may not be the next sequential request.
- 1 (Nearest) Selects request that is closest to the current cylinder.

2 (Elevator) - Processes requests as it moves in one direction along the disk until it reaches the last request in that direction. This means that favors the center of the disk, as it passes it twice as often as the periphery.

3 (Forward) - This processes requests from the lowest cylinder number to the highest in direction only. Note that optimization is effective if the host operating system supports multiple accesses.

Fairness count:

The fairness count determines the number of times an I/O request will be passed over by the controller's seek optimization setting before it is executed. A reasonable count for normal use would be around 25. Every time a request is passed over, its fairness count is decremented. When that count reaches zero that request will be selected, no matter what optimization strategy is in effect. This count has no effect if no optimization is selected.

This completes the drive structure specification. WOMBAT next prompts for a unit number. Type the unit number of the next unit to be defined, then hit RETURN. When the unit is completely defined, WOMBAT will again prompt for a unit number. If there are no more units to define hit RETURN. WOMBAT will then proceed to the controller wide parameter definition. Note that if no units are defined, the operating system will not see anything attached to the controller.

Unit Structure Specification

Unit size:

If an existing unit number is specified WOMBAT displays the size in blocks. If a new unit is specified WOMBAT displays the size in of the first unallocated disk area it finds at the start of the disk. On a new disk this will be the entire user area. This can be changed to smaller value if necessary. To delete an existing unit, specify zero for this field.

Media type:

This field is displayed by some operating systems when you inquire about the type of drive. As a part of unit status when a "Get Unit Status" command is issued the MSCP protocol returns a 5-character media type. The first two characters must be 'DU'- for example - DURD54. To change this enter 1 to alphabetic characters and 2 digits, e.g. to emulate DEC's megabyte Winchester.

Serial number:

The MSCP protocol returns a 32-bit volume serial number as a part of its response when an "on-line" command is issued. WOMBAT defaults this field to zero. To change this enter the desired serial number.

Host write confirmation:

This specifies if the controller is to notify the host that a write request has been completed. Enter 1 if the host is to be notified when the data is in the cache. The data will be written to the disk later. Enter 2 if the host is to be notified when data has actually been written to the disk.

This completes the unit structure definition. Now WOMBAT enters the final part of the dialogue which specifies the controller wide tuning parameters.

Controller Wide Tuning Parameters

Read lookahead:

This is a feature of the cache which allows the controller to read a specified number of sectors in addition to those requested by the host. Enter the minimum number of blocks you wish controller to read for any request. For example, if a value of 4 is specified, when the host asks for a single block to be read, the next 3 will automatically be read into the cache. If the host subsequently asks for one of these

blocks then the request can be honored from cache. If the host requests a transfer to or larger than the read lookahead then this parameter will have no effect.

Cache Enable:

This parameter enables or disables the cache. A value of 0 disables the cache and a value of 1 enables cache. The cache cannot be selectively enabled or disabled for a particular drive or unit number.

Command queue:

This parameter allows you to specify commands the controller can stack. The controller will then attempt to optimize the order in which they are executed. Large command queue stacks incur considerable overhead and will degrade controller performance. The size [8.] is a good compromise.

Front Panel Type:

This option allows the correct selection of front type. See Front Panel Connections.

0 - None 1 - Passive 2 - Active

This completes the disk structure definition. WOMBAT now checks the tables for consistency and returns to the disk structure menu.

Option 2 - Format Disk

WOMBAT asks you to confirm this drastic action as it will destroy ALL data that resides on the disk. WOMBAT will then initiate a two pass formatting operation. During the first pass WOMBAT creates all the sectors on the disk, writes sector headers which contain the sector number, the head number, and the cylinder number, as well as preambles and sync bytes, followed by a 2 byte data field. During the second pass WOMBAT writes a test pattern to each sector, preparing the disk for read testing. WOMBAT then writes the disk structure onto the reserved areas.

Option 3 - Write Disk Structure

WOMBAT will ask you to confirm this drastic action as any existing disk structure will be destroyed. WOMBAT will then write the new structure onto special reserved areas of track zero. The data is recorded twice for improved recoverability. A total of 6 blocks is written on track zero. After the structure has been written, the drive's replacement block table is zeroed. If there were any replaced blocks recorded there, they will be lost. However, they will still be marked as replaced and will generate hard errors during a read operation.

Option 4 - Update Header Blocks

This is similar to Write disk structure except that the replacement block table is not written, thus preserving any blocks which may have been replaced. This option is used after changes to the disk structure such as changing unit numbers or redefining the virtual units. Unless the header blocks have been updated, the changes are not recorded on disk.

Option 5 - Display Disk Structure

WOMBAT displays the structure of the currently selected drive in a form similar to the create disk structure dialogue. This is useful for checking that the newly created structure is correct.

Option 6 - Change Unit Number

It is sometimes necessary to change a unit number in order to resolve a duplicate unit number or to satisfy operating system requirements. This method is a safe and simple way of doing so. WOMBAT prompts for a unit number on the current drive, and then for the new number for that unit. NOTE: For the change to take effect, the header blocks must be updated using option 4 above.

Display Statistics

Statistics about disk and cache usage are maintained, and recorded on the disk periodically. They are displayed as:

		Controller statistics report
#	of commands	XXXX
#	of reads	xxxx # of cache hits xxxx (xxx%)
#	of writes	XXXX

Drive statistics report

Drive # Soft errors Re-vectors Blocks replaced xxx xxxx xxxx

Drive # Seek distance # of seeks Seek errors xxx xxxx xxxx xxxx

Commands is the number of MSCP commands issued.

Re-vectors is the number of accesses to replaced blocks.

Seek distance is the total seek distance, in cylinders.

Blocks replaced is the number of blocks dynamically replaced by the controller during normal operation, rather than through WOMBAT.

Reset Counters is then asked. "Y" will reset them to zero.

6.4.6 Disk Test Menu Options

A disk can be tested after it has been formatted and before the structure is written to it. Testing does not overwrite the HDR or RCT blocks. The disk structure must be written to the disk before bad blocks can be replaced.

All tests continue indefinitely until aborted by one of the following methods:

- 1. If an RS232 serial port terminal is attached to the controller, press BREAK.
- If WOMBAT is running from the Console terminal, type CTRL-C.

When a test is aborted the Test Disk Menu options are returned. If tests are run from an RS232 terminal attached to the controller, beware of system activity on the host computer as Qbus initializations will cause the disk controller firmware to re-initialize and so leave WOMBAT.

All tests give 10 retries on an error, reporting every error by displaying the block number and an error code.

- ** Disk Test Menu **
- (! means all data on Disk destroyed)
- 1 Read All Disk (preserves all data)
- 2 Write Disk

- 3 Pattern Test
- 4 Random Write
- 5 ECC Validation
- 6 Read Physical Block
- 7 Display Error Statistic
- 8 Zero Error Statistics
- 9 Test Cache, RAM & ROM

Option 1 - Read All Disk

This test reports any read errors. Successful operation will be reported in the following format:

Pass: 1. Errors: 0. Pass: 2. Errors: 0.

This function does not destroy any information.

Option 2 - Write Disk

This test reports any write errors while writing a test pattern to the whole disk. ALL INFORMATION on the disk, excepting HDR and RCT blocks, is DESTROYED. Errors are displayed in the standard format:

Block: 32040 (Error message)

Pass: Errors: 1
Pass: Errors: 1

The displayed error count is cumulative until the test is terminated.

Option 3 - Pattern Test

A test pattern is written to each block, including the replacement blocks. WOMBAT does one write and 10 read passes. This test reports any errors in the standard format as shown above.

Option 4 - Random Writes

This test writes 5000 blocks at random locations in the user area of the disk. It then reads the entire disk to determine if any of the writes caused an error. This test is designed to test the head positioning and selecting logic of the drive.

Option 5 - ECC Validation

The ECC test uses a special reserved block on track zero for testing. It first proves that it can successfully correct an 11 bit error and then proves that it cannot correct a 12 bit error. This test checks the ECC logic within the QDA4E.

Option 6 - Read Physical Block

WOMBAT prompts for a block number anywhere on the disk. It then converts that block number into a physical address consisting of cylinder, head, and sector, and displays these values in hex and ASCII. Then it reads that sector and displays a message indicating the success or failure of the read. The same physical block can be re-read by typing \ instead of a block number. The block's replacement block can be read by typing @.

Option 7 - Display Error Statistics

Displays the error statistics gathered by any of the above disk testing options in the following format:

** Error Statistics **
Block Number (of errors)
32040 1.
Blocks in error: 1.

Option 8 - Zero Error Statistics

Zeroes the error statistics table & redisplays Test Menu options.

Option 9 - Test Cache, RAM & ROM Option

This test continuously writes test patterns throughout the entire cache and reads them back testing for veracity. A separate part of the test automatically checks that the parity logic is functioning correctly by forcing incorrect parity and checking that an error occurred. The cache pattern tests use special microcode instructions which iteratively read and write large blocks of cache memory at high speed. The Static RAM is also tested and the code PROM is checksummed.

6.4.7 Bad Block Management Menu Options ** Bad Block Management Menu **

- 1 Manually Replace Bad Block
- 2 Automatically Replace Bad Blocks
- 3 Display Replaced Bad Blocks
- 4 Enter defect map
- 5 Get defect map from drive

Option 1 - Manually Replace Bad Blocks

WOMBAT prompts for a block number within the user area of the disk. Then it marks the specified block as bad and allocates a replacement block for it.

Option 2 - Automatically Replace Bad Blocks

WOMBAT searches the error statistics table, which is compiled by the read, write, and pattern tests, for blocks whose error count exceeds three. Any such blocks are marked as bad on the disk and replacement blocks are allocated for them.

Option 3 - Display Replaced Bad Blocks

WOMBAT reads the Replacement Control Table and displays the logical block numbers of any blocks recorded there.

Option 4 - Enter Defect Map

This option prompts for the drive manufacturer's defect map information as follows:

Enter defect map for drive n

Enter all values in decimal

Bytes per sector: The number of bytes per sector as set up in the drive switches (hard sector drives), or calculated by following the procedure in the Computing Sectors per Track section.

Cylinder: Cylinder number of defect, RETURN or

CTRL-C if no more.

Head: Head number of defect.

Bytes past index: Location past index of defect.

Bit length of defect: Length of defect in bits.

This data is then used to compute the address of a block on the disk. If it does not match a block on the specified disk track the error message "!! Beyond last sector" is produced. This may mean that the defect is located beyond the last data sector on the track, or that the entered data was wrong. WOMBAT then calculates the block number, displays it and replaces it. The "Cylinder:" prompt is then repeated.

Option 5 - Get Defect Map from Drive

This option tells the controller to find then read the defect map information written on the drive by the manufacturer. It will then list the defective blocks by head number and ask for confirmation of replacement. It is important that the displayed data be checked, as it may have been corrupted by previous formatting/overwriting. WOMBAT does do some checks, but relies partially on the operator checking the validity of the list, in particular, the fields underlined below.

There should be a printed copy of the defect map supplied with the drive and this should be used for checking.

ESDI defect map for drive n

Soft/Hard sectored drive, xxx bytes/sector

Defect map read error, head x

Displayed if the defect map is unusable.

Then, for each head on the drive:

Defect map read: Date - mm-dd-yy Fujitsu mm-dd-yy)

NOTE

The date code in brackets (Fujitsu mm-dd-yy) is displayed because early model Fujitsu drives recorded the date incorrectly.

The date code should match the date on the printed Defect List supplied with the drive.

Head x Defect list for head x

Cyl. xxx, bytes past index xxx, length xxx

This is displayed for each error.

The Defect Map for each head has the head number recorded. If WOMBAT detects a mismatch it reports it and asks if you wish to proceed when all the defects have been displayed,

Do you wish to proceed with replacement for head x?

Only reply "Y" if the list appears correct. A defect at cylinder 200, 30,000 bytes past the index, and of length 0 is a possible (but patently wrong) display.

The defects are again displayed as they are replaced:

Cyl. xxx, bytes past ind. xxx, length xxx: block xxx, byte xxx

When this is complete, the message

Defect map replacement complete for head x

is displayed, and the data for the next head processed.

6.4.8 Shadow Options Menu Option

This feature shadows, i.e. keeps two copies, of a 'logical unit'. When the disk controller and drives are first powered up, any unit (and a physical disk may be broken up into a number of 'logical units') that is shadowed has its contents copied to its shadow unit. Thereafter any update of that unit will cause the controller to update the shadow unit as well. When reading from a shadowed unit, the drive with its heads nearest the required data is used. This helps to keep the drive shadowing overhead down, although in normal circumstances, with writes consisting about 10% of reads, there is a performance penalty.

Shadowing consists of two operations: the initial copy of the entire primary unit to the shadow unit; and then the updating of both primary and shadow units on every disk write.

Updating of both copies takes place as a matter of course, while the initial copying can happen at various times depending on the 'Shadow copy option' selected.

The main reason for drive shadowing is RELIABILITY. When reading, any error detected will cause the controller to use the data on the other drive. This means that, besides the controller having extensive error recovery facilities, there is a redundant backup of all data on the shadowed unit without any user programming or operating system overhead.

Drive shadowing is completely controlled from the 'Shadow options' menu.

- ** Shadow Options Menu **
- 1 Shadow copy option
- 2 Set shadow units
- 3 Reset shadow units
- 4 Copy unit to unit
- 5 Compare unit to unit

Option 1 - Shadow Copy Options

The first step is to select how the unit (the 'Primary unit') is to be copied to its 'Shadow unit'. There are 4 options:

- O WOMBAT Only

 This is the default. With this selected the primary is only copied to its shadow when you invoke it in the Copy Unit to Unit option.
- 1 On power up

 This causes the copying to be done whenever the drive(s) are first powered up, or WOMBAT has been invoked.
- 2 Not ready/ready The copying is to be done whenever the controller detects the shadow unit going from a not ready to ready condition e.g. a removable drive being brought on-line.

The use of copy type 0 only is recommended at this time. It is considered that the other options have too high a

probability of copying bad data from the primary unit over good backup data on the shadow unit.

Option 2 - Set Shadow Units

Shadow units are logically connected to primary units by this option which asks:

Primary unit: The unit to be shadowed.

Shadow unit: The unit number of the shadow.

The shadow unit will not be available to any operating system as it appears as 'unit undefined'. It must be EXACTLY the same size as the primary unit.

Now copy the primary to its shadow by using the Copy unit to unit option.

You may shadow a unit on the same physical drive, though you lose both reliability, by having the data on only one drive, and performance, by having the disk heads move to write all data twice.

If you have selected a shadow copy type of 1 or 3, whenever the controller is powered up and its first 'MSCP Initialize' sequence executed for the primary unit (the unit you wanted shadowed), the contents of this unit are copied to its shadow, or if WOMBAT has been invoked. The time this takes obviously depends upon the number of blocks on the unit, but we have found that a third of a megabyte per second is a good guide. While this is happening user I/O may take place, but will be made slower by the 'shadow copy' taking place.

WARNING

A potentially serious problem exists with the automatic copy of data from the primary to the shadow at power up (shadow copy types 1 or 3). If the primary has been corrupted, or the data is invalid, IT WILL DESTROY YOUR BACKUP. If you have trouble with your primary unit use WOMBAT to change it from a shadow primary. You may even want to change its unit number via the 'Change Unit number' option, and then bring in the shadow unit as its replacement. Please realize that the controller MUST assume the validity of the primary unit at the first initialization after power up if you have selected one of the power up copy types.

Option 3 - Reset shadow units

This option allows disabling of an assigned shadow unit:

Primary unit: The primary unit.

Shadow unit: The unit number of the shadow to be

disconnected from the primary.

This shadow unit will now be available to the operating system.

Option 4 - Copy Unit To Unit

To copy a primary unit to its shadow:

Primary unit The number of the unit you want

copied.

Shadow unit The unit you want the data copied

to.

The copy will take place. Data is transferred at about twenty megabytes a minute

Option 5 - Compare Unit to Unit

This option allows the comparison of the data on a primary and a shadow unit.

Primary unit One of the units.

Shadow unit The other.

If any data errors, or data compare mismatch, is found, the block number and type of error will be reported. The comparison proceeds at about two megabytes a minute.

WARNING

When the primary drive fails, the QDA4E does not automatically switch to the secondary drive. If the primary drive fails you must shutdown system and reboot: go into WOMBAT and select option 5, then option 3, and reset the secondary drive back to normal. Exit WOMBAT and reboot system--which should boot from the previously shadowed drive.

6.4.9 WOMBAT Disk Structure

WOMBAT records the logical structure of the drive on track 0. All of track 0 is reserved for this and other testing purposes. The user area begins at the next block after track

O. This is the same as the number of sectors per track. The user area extends to the block before the beginning of the Replacement Control Table (RCT). The size of the RCT is fixed and is always 2 tracks. The 2 tracks are accessed by different heads and contain identical data. All the blocks from the end of the second RCT track to the end of the media are reserved for replacement blocks.

The disk area reserved for RCT and replacement blocks is always an integral number of cylinders. The number of cylinders required is always computed from the total number of cylinders on the drive. This reserves a constant proportion of the media for replacement blocks. The amount reserved is approximately 0.1% of the total formatted capacity. Option 4 of the master menu will report the formatted capacity of each drive. You can determine the user area easily. It will be reported as the size of the first unit configured on that drive. If there is more than 1 unit, then the sum of their sizes is the total user area, assuming there are no unallocated areas on the disk.

All block numbers in WOMBAT are physical block numbers beginning at the first sector of the first head of the first cylinder of the drive, which is defined as block zero. The last block on the drive is block n-1 where n is the total number of blocks on the drive. Therefore, the first block of the user area is not block zero. Its block number is the same as the number of sectors per track, as track 0 on the drive is reserved. Option 4 of the master menu will display both the size and the offset (starting block number) of each unit defined. Using these figures you can determine the exact position and extent of any unit.

6.4.10 Drives with Removable Media

Some manufacturers offer drives with removable media. The QDA4E detects when such a drive is connected and makes certain changes to the way it operates. The important changes are that you can only have one virtual unit on a drive whose media can be removed. That unit must have the same unit number as the physical drive number. The reason for this is that it is necessary to determine the unit number even though the media may be removed. This means that the unit number cannot be recorded on the media for these drives. WOMBAT automatically detects these drives and prevents you from creating incorrect units on them.

6.4.11 Error Recovery Procedures

During normal operation, the QDA4E checks every disk transfer for errors. When an I/O error is detected, the QDA4E enters a special error recovery procedure to attempt to provide the host with 'perfect' media.

The first method used to recover the data is simply to try the operation again. If this succeeds, the host is guaranteed to receive good data. This is repeated until a threshold is reached, at which time the second error recovery procedure is initiated. The second procedure recalibrates the drive and reseeks to the block in error in an attempt to correct any positioner errors that may have prevented data recovery. After the reseek, the operation is again retried until the retry threshold is reached. If any of these retries succeeds then the host is guaranteed good data. If the retry threshold is reached after reseeking, then the final error recovery procedure is attempted. This procedure is ECC. The ECC bits appended to each block are used in an attempt to correct the error. If successful, the host is guaranteed good data.

If the data is successfully recovered by retrying, then the number of retries necessary is checked. If it exceeds the retry soft error limit then that block is dynamically replaced, on the assumption that it is in the process of gradual failure and will get worse. The known good data is written to the replacement block and the host is notified of success.

If the data is successfully recovered by the ECC algorithm then the block is dynamically replaced on the assumption that the block has developed a hard error. The corrected data is written to the replacement block and the host is notified of success.

If the data cannot be recovered by any of these means then the block is assumed to be bad. The block is dynamically replaced, and is written with forced error status. This will cause forced error status to be returned whenever the block is read, telling the host that while the block is good, the data in the block is bad. When the block is written, the forced error status will be removed and from then on the block will be good. When a block is dynamically replaced with forced error status, the host is notified of a forced error on that block.

6.4.12 Drive shadowing

The intent of this feature is to provide a continuous automatic backup of important data. The normal case would be to have two identical drives with a single unit on each. The primary unit must be of identical size to the shadow. Once the primary/shadow pair has been defined, then the shadow unit will become invisible to the host. All writes directed to the primary will also be written to the shadow. If a read is directed to the primary, the controller will attempt to redirect that read to whichever drive of the pair has its heads positioned closest to the requested data.

This feature will only offer protection against the primary unit failing. It cannot protect, as a normal backup can, against a rogue program which writes garbage on the disk. The garbage will simply be copied onto the shadow unit as well as duplicating the garbage. It is advisable, although not mandatory, that the shadow unit be on a different drive to the primary if the maximum protection is to be gained.

If the primary does fail, then you will want to recover the data from the shadow. To do this WOMBAT must be run. First, disable the shadow unit with the "Reset shadow units" option of the Shadow Options Menu. Then (optionally) change the unit numbers of the two units with the "Change unit number" option of the Disk Structure Menu. Once the data has been recovered, the primary unit can be serviced and re-installed in the system.

6.4.13 Computing Sectors Per Track

To compute the number of sectors per track first determine the number of bytes per track from the drive manual. Divide this figure by the number of bytes per sector. You determine the number of bytes per sector by the following calculation:

Bytes per sector = K + (PLO * 2)

WHERE: K = 566 for hard sectored drives K = 557 for soft sectored drives

PLO is the number of bytes of PLO sync required by the drive. Obtain this figure from the drive manual.

e.g. Maxtor EXT4125, PLO sync field is 26 bytes bytes/sector = 557 + (26 * 2) = 609 bytes per sector

This drive has 20808 bytes/track minimum sectors/track = 20808/609= 34.17 so choose 34 sectors per track

e.g. Hitachi DK512 PLO sync field is 11 bytes bytes/sector = 566 + (11 * 2) = 588 bytes per sector

This drive has 20832 bytes/track minimum sectors/track = 20832 / 588 = 35.42 so choose 35 sectors per track

6.4.14 WOMBAT Error Messages

The following is a list of the error messages displayed by WOMBAT.

sector not found	Indicates that the sector asked does not exist or cannot be located.
drive fault	Indicates that the drive is faulty and that service is necessary.
drive timed out	Indicates that the drive has failed to complete an operation.
data field error	Indicates that bad data exists in a sector.
controller fault	Indicates controller failure. Service will be necessary.
block marked as bad	Indicates that the block has been flagged as bad. The controller will refer to the RCT for a replacement block.
data late	Data is lost due to internal overflow

in controller memory before

transmission over the bus to host.

QDA4E CONTROLLER

QDA4E SUBSYSTEM USER'S GUIDE

forced error	A forced error occurs when a good block has bad data. The block is flagged forced error until good data is written to the block.
seek error	An error has occurred on a seek operation.
rct full	The replacement control table is full. An error of this kind indicates that the disk has too many errors to be serviceable.
rct read error	Fatal error which indicates the controller cannot read the RCT.
rct write error	Fatal error which indicates the controller cannot write the location of replacement blocks.
illegal sector specified	A sector with an illegal number has been specified.
illegal block number	A block with an illegal number has been specified.
non existent drive	A drive has been specified which does not exist.
unit table full	The unit table is full. The maximum number of units is 16. To create a new

nonexistent unit The unit selected has not been defined.

unit an existing unit must be undefined.

A drive must be selected before WOMBAT can perform any operation. Select any valid drive.

no drive selected

disk structure write error

WOMBAT is unable to write the structure to the disk. This is a fatal error indicating that the drive is not serviceable.

6.4.15 WOMBAT Self-Diagnostics Initialization Procedures

A common initialization procedure exists for both WOMBAT and the MSCP firmware. It performs:

a RAM integrity test a ROM checksum various checks on the disk drive and its structure

The errors which can result from this are described under Fatal Controller Errors section.

6.5 MSCP Programming

6.5.1 Overview of MSCP

Mass storage control protocol (MSCP) is a message oriented set of rules by which the QDA4E controller module and the host system communicate. This protocol allows the host to send message requests for data reads or writes to the controller and received response messages back from the controller. The host does not concern itself with details such as device type, media geometry, media format, or error recovery.

All software and hardware functions are partitioned into independent 'host' and 'controller' layers. Each layer consists of a high-level I/O driver and a communications server. The controller layer receives and processes commands which have been formed by the host layer.

The communications server handles all communications protocol between the I/O layers, leaving the I/O system free to process data requests. Communications between host and controller are carried out on the I/O bus without having to generate processor interrupts. The host's communications server monitors all command transmission and response and in the event of failure or error, initiates recovery procedures.

Disk drive parameters are transparent to both the host and controller resident layers of MSCP. The disk drive passes factors such as sick geometry, storage capacity or error retry counts to the disk controller on system start-up.

In addition to relieving the host of disk-specific data, the disk controller and disk provide the host with "clean" data. The disk drive handles some positioner errors entirely by itself but performs certain error-recovery operations under direction of the disk controller.

6.5.2 Controller Communications

The host designates an area of memory to be used as a communications area between itself and the controller. This area is made up of two sections, a header area containing interrupt identification words and a variable-length section containing the response (receive) and command (send) lists, organized into ring buffers. Following is the memory communications format.

MEMORY COMMUNICATIONS FORMAT

	<>	
Lower Address	! COMMAND INTERRUPT WORD !	
Address	! RESPONSE INTERRUPT WORD !	
	! TWO WORD !) ! BUFFER DESCRIPTOR !) !!)	Response
	!! !!	descriptor ring
Two MSBs = O	!! !! !! !!) Command) descriptor) ring
	· · · · · · · · · · · · · · · · · · ·	{
Highest	!!)
Address	!!)

Command and Response Rings

Command and response lists are organized into 'rings' of 32 bit descriptors. The length of each ring is determined by

the speed with which the host and controller generate and process messages. The host sets the right lengths at initialization time. The following describes the Descriptor Format while table below details the code descriptions.

DESCRIPTOR FORMAT

15	5 				 																							C)0
					 	L	!	L	!		!	L	!	L	!	L	!	L	!	L	!	L	!				!	Z	!
!	0	!	F	!	 					ese														H	!	Н	!	H	!

COMMAND RING CODE DESCRIPTIONS

CODE	DESCRIPTION
z	Is zero, as envelope address (text+0) is word- aligned. The controller will always assume that Bit 00 is set to zero.
L	Low-order envelope address.
Н	High-order envelope address.
F	Flag bit.

When the controller returns ownership to the host it sets F=1 to indicate that it has completed action on the descriptor.

When the controller acquires ownership of a descriptor from the host, F=1 indicates that the host is requesting a ring transition interrupt. If F=0, the host is not requesting a ring transition interrupt. An interrupt will occur only if this descriptor causes a ring transition and if transition interrupts were enabled during initialization.

The controller always sets F=1 when returning a descriptor to the host, so if a host wishes to override ring transition interrupts it must always clear F when passing ownership of a descriptor to the controller.

O Ownership bit. Set to O if owned by the host or 1 if owned by the controller. Interlocks the descriptor against premature access by either party.

Message Packets

The command or response descriptor points to word (text+0) of a 16-bit word-aligned message envelope formatted as shown below. Table below describes the word envelope contents.

MESSAGE ENVELOPE FORMAT

	15		80	07	04	03	00		
-4	!	Message	lengt	h (:	in byt	es)	!)	Message envelope
-2	!	Connection	Id!	ms	sgtyp	! c	redits!)	enverope
text+0	!	MB1	!		M	IBO	!		
+2	!	MB3	!		IM.	IB2	!		
	•						•		
	•						•	•	•
	!	MBn-1	!		M	íBn-	2 !		

WORD ENVELOPE CONTENTS

WORD	ENVELOPE CONTENTS
0	Message length, in bytes.
	For commands, this length is equal to the size of the command (in bytes), beginning with (text+0).
	For responses, the host sets the length equal to the size of the response buffer (in bytes), beginning with (text+0). Before actual transmission of a response, the controller reads the

field length in the message envelope. If the controller's response is longer than the response buffer, the controller will fragment its response into as many response buffers as necessary.

The controller sets the resulting value into the message length field. The host must therefore keep re-initializing the value of this field for each proposed response. If a controller's responses are less than or equal to 60 bytes, then the controller need not check the size of the response slot.

1 Connection Id

Identifies the connection serving as a source of, or destination for, the message in question.

2 Message Type

The following response ring message types are implemented:

MSGMNT Maintenance packet (diagnostic)
MSGCRD Credit notice (ignored)
MSGDAT Datagram packet
MSGSEQ Sequential packet

3 Credit Field

Gives a credit value (usually one) associated with the message. This mask, in response packets, is added to the controller's credit field to give the number of commands-in-progress. So while Word 1 is always 1 for the command ring, this is not the case for response rings.

6.5.3 Message Transmission

Command Transmission

When the ownership bit (0) of a command ring descriptor is equal to 1, it means that the host has filled the descriptor and is releasing it to the controller. When the ownership bit (0) resets to zero, it means that the controller has emptied the command ring descriptor and is returning ownership of the descriptor to the host.

To ensure that the controller sees every command, the host must read the IP register whenever it inserts a command in the command ring. This forces the controller to poll the command if it was not already accessing the command ring.

Response Transmission

This forces the controller to poll the command if it was not already accessing the command ring.

Response Transmission

When the ownership bit (0) of a response ring descriptor is equal to zero, it means that the controller has filled the descriptor and is releasing it to the host. When the ownership bit (0) sets to 1 it means that the host has emptied the response ring descriptor and is returning ownership of the descriptor to the controller. Just as the controller must poll for commands, so must the host poll for responses.

Interrupts

The transmission of a message will result in a host interrupt from the controller under the following circumstances.

- 1. During the initialization process (open a 'connection').
- 2. When the command ring buffer transitions from 'full' to 'not empty'. This interrupt means that there is a response for the host to process.
- 4. When a fatal controller error is detected and an interrupt can be generated. These are:

Failure to become Qbus master for data transfer Failure to become Qbus master for interrupt Failure to access I/) page registers or communication area

Qbus parity error detected.

6.5.4 Data Transmission

In the command ring, the descriptor points to a command packet. Within the command packet is a buffer descriptor which contains a pointer and a byte or word count. The buffer descriptor points to the data buffer which holds data transfers. The data is moved by the controller into or out of the buffer as DMA transfers to/from Qbus addresses.

6.5.5 Initialization

The purpose of initialization is to identify the parameters of the host-resident communications region to the controller, provide a confidence check of controller integrity, and bring the controller online to the host.

Initialization Process

This paragraph describes the activity within the SA register during an initialization process. This is dependent on whether SA is being read or written.

By moving 4000 into IP, the controller initializes and passes back the 'step' response in SA. Then, the initialization parameters are written into SA. There are 4 words of initialization, and the controller must reflect each step by the appropriate step response, which is also returned in SA. The initialization parameters are in the following table.

INITIALIZATION PARAMETERS

WORD	CONTENTS
0	Command and Response ring sizes, interrupt enable and vector. The host writes into SA the lengths of the rings, whether interrupts are to be armed, and if so, the address of the interrupt vector. The controller then runs a complete internal integrity check and signals either success or failure.

- Low order address of communications area, ie., ring buffer address. The host reads an echo of the ring lengths from SA, and then writes into SA the low-order portion of the ring base address.
- High order address of communications area, bits 0-14. The interrupt vector address and the master interrupt arming signal are echoed in SA. The host then writes the high order portion of the ring base address to SA along with a signal that conditionally triggers an immediate test of the polling functions of the controller.
- Burst transfer control, last failure flag, and the 'GO' bit. The controller tests the ability of the Qbus to perform DMA transfers. If successful, the controller zeros the entire communications area, and then signals the host that initialization is complete.

6.5.6 Registers

The programmable registers contained on the QDA4E are the Initialize and Poll register (IP) and the Status and Address register (SA).

Initialize and Poll Register (IP)

The host begins the initialization sequence by either issuing a bus initialize or by using the IP initialize operation. Any write to that address will cause an initialization of the controller. When read while the controller is operating, it causes the controller to initiate polling. The QDA4E responds to the 16 bit initialization words as set out in following table.

ODA4E	TNTTI	TITA	ZATION	WORDS

NUMBI OCTAL	ER HEX	FUNCTION	PROCESSOR
000254	00AC	WOMBAT	MicroVAX II On board maintenance port User communication program
000260	00B0	WOMBAT	
000272	00BA	WOMBAT	

Status and Address Register (SA)

The SA register consists of a set of two registers, the SA read register and the SA write register.

When read by the host during initialization, it communicates data and error information relating to the initialization process. When written by the host during initialization, it communicates certain host-specific parameters to the controller.

When read by the host during normal operation, it communicates status information including fatal errors detected by the controller.

6.5.7 MSCP Commands

QDA4E MSCP COMMANDS

COMMAND	FUNCTION
Access	Reads data from the specified unit.
Abort	Guarantees that referenced MSCP command will complete within the controller timeout period.
Available	If specified unit is on-line, returns it to the unit-available state. If specified unit is currently in the unit-available state, this command has no affect.
Compare Host Data	Reads data from the disk and compares it with the data in the host buffer.
Erase	Writes zeros to the specified logical blocks on the unit. (No data is accessed from the host).

Get Command Status	Reports on the status of a specified command by returning a number that reflects the command's progress.
Get Unit Status	Reports on the status of a specified unit.
On Line	Places the specified unit on line, if possible.
Read	Reads data starting from the specified logical block on the disk, into host memory.
Set Controller Characteristics	Sets host-settable controller characteristics.
Set Unit Characteristics	Sets host-settable unit characteristics.
Write	Writes data starting at the specified logical block on the disk, from the host memory.

6.5.8 Error Handling

High data integrity is achieved by the controller through a 48 bit ECC (error checking and correction) polynomial with an 11 bit correction span. ECC will first try to read or write a block up to 10 times before attempting to correct the error. If error correction fails a non-recoverable error is reported.

6.5.9 Fatal Controller Error

If a fatal error is detected when the controller is uninitialized, the red error LED flashes and the fatal error status is set in the SA register. Following table describes fatal controller errors.

FATAL CONTROLLER ERRORS

ERR octal	OR hex	DESCRIPTION
100004	8004	RAM test failure
100005	8005	ROM checksum failure
100011	8009	No drive
100100	8040	Disk unformatted
100101	8041	Disk unstructured

Running WOMBAT and selecting the Display Error option will give an appropriate error message. A full description may be found under WOMBAT Error Messages section.

MSCP STATUS CODE MESSAGES

MESSAGE	MEANING
Command Aborted	The current command was aborted before it could be completed normally.
Compare Error	While performing a Compare command, a Discrepancy was found while comparing the disk data to the host data.
Controller Error	The QDA4E controller detected an internal error, but is able to continue processing its outstanding commands.
Data Error	An error was detected in the reading or writing of data. ECC attempts to read

or write data up to 10 times. If the error persists correction is attempted. If correction fails the error is reported.

Drive Error A drive-related error was detected (such

as a seek failure).

Media Format Error Indicates that the media mounted on the

unit was incorrectly formatted.

Host Buffer Access Reports bus timeouts and parity errors Error during data transfers. (Applies only to

during data transfers. (Applies only to the data portion of an MSCP command).

Invalid Command The ODA4E controller found some field in

the command to be in error.

Success The command was successfully completed.

Unit Available The QDA4E controller is not on line, but

it can accept an On Line command from

the host.

Unit Offline The QDA4E controller is not on line, and

it cannot be brought on line.

Write Protected A Write or Erase command was attempted

to a unit that is logically

write-protected.

6.6 QDA4E Operating Systems

The following discussion is intended to supplement DEC operating system resources and aims to aid the user of the QDA4E in understanding how different operating systems integrate the device. This information will help the user of the controller plan the installation and in choosing the appropriate bus addresses and interrupt vectors for the disk subsystem. For a complete description the DEC system documentation should be consulted.

6.6.1 Operating Systems Overview

In order to install any new device in a computer, the host operating system must be informed of the device's existence and where to find that device. In DEC operating systems this can be done in one of the following ways:

- (a) The device can be manually connected using Connect or Configure statements.
- (b) The operating system can be informed about the peripheral device during an interactive SYSGEN.
- (c) The operating system can poll the device I/O address space.

Any of these methods will accomplish the desired result. The host system will be alerted to the device's existence, type, address and interrupt vectors.

Method (a) creates a command file that is executed on power-up. Method (b), interactive SYSGEN creates a configuration file that the operating system accesses on power-up. Method (d) is referred to as 'autoconfigure'.

MSCP Devices

The QDA4E is an MSCP (Mass Storage Control Protocol) type device. All MSCP-type devices contain two registers that are visible to the Qbus I/O page. They are the Initialization and Polling (IP) register and the Status and Address (SA) register.

Obus Addresses

The standard Qbus address of 17772150 (Octal) is used as the address of the first controller on the host system. The IP register, CSR address, Qbus address, and the base address all refer to the same register.

Vector Addresses

Many operating systems choose vector addresses automatically. If an operating system requires manual input of vector addresses they are programmed into the controller during the

initialization process.

Device Names

Following table lists device names for MicroVAX II. Two controller and device names are given to indicate the numbering scheme.

DEVICE NAMES IN DEC OPERATING SYSTEMS

VAX/VMS PUA, PUB DUAO, DUA1	OPERATING		OLLER	DRI	
VAX/VMS PUA, PUB DUAO, DUA1		1st:	2nd:	1st:	2na:
	VAX/VMS	PUA,	PUB	DUAO,	DUA1

6.6.2 MicroVMS Operating Systems

The first QDA4E controller is located at the standard bus address of 1777 2150 (Octal) and the second in floating address space. The MicroVMS SYSGEN utility can determine the Qbus and interrupt vector addresses for any of the I/O devices installed on the bus. MicroVAX/MicroVMS must be running in order to use this utility. The Qbus and interrupt vector addresses can be determined manually if access to a running system is not possible.

Using MicroVAX/MicroVMS SYSGEN

The following is an outline of the MicroVMS SYSGEN procedure to determine Qbus and Interrupt vector addresses. This procedure requires system manager privileges.

1. Login and run the SYSGEN utility.

\$ RUN SYS\$SYSTEM:SYSGEN<return> SYSGEN>

The SYSGEN> prompt indicates that the program is ready.

2. Obtain a list of the devices currently installed on the MicroVAX Qbus by typing:

SYSGEN> SHOW/CONFIGURATION<return>

and get :-

Name: PUA Units: 1 Nexus: 0 CSR: 772150 Vector1: 154

Vector2: 000

Name: TXA Units: 1 Nexus: 0 CSR: 760500*Vector1:

310*Vector2: 000

* Indicates a floating vector or address.

Sysgen lists the devices already installed on the Qbus by logical name. Devices with floating bus and vector addresses should be noted if it is intended to re-install them with the QDA4E controller. Floating bus addresses will be larger than 760000 (octal). Floating interrupt vectors will be larger than 760000 (octal). Floating interrupt vectors will be larger than 300 (octal).

3. Execute the configure command. This will determine the Qbus and Vector addresses that autoconfigure will expect for each device type.

SYSGEN> CONFIGURE<return> DEVICE>

Specify the devices to be installed on the bus by typing their Qbus names. Under MicroVAX/MicroVMS the device name for MSCP-type controllers is UDA.

DEVICE> UDA,2<return>
DEVICE> DHV11<return>

The device name is separated from the number of devices by a comma. The number of devices is specified in decimal. Devices with floating addresses or vectors are not affected by devices with fixed addresses or vectors. Only devices with floating addresses or vectors need be specified.

 When all the devices have been specified enter a control-Z.

DEVICE> CTRL-Z

The addresses and vectors of the devices entered will be listed in the following manner:

Device: UDA Name: PUA CSR: 772150 Vector: 154 Support: yes Device: UDA Name: PUB CSR: 760334 Vector: 300 Support: yes Device: DHV11 Name:TXA CSR: 760500 Vector: 310 Support: yes

- * Denotes floating bus and interrupt vector addresses. Floating CSR addresses must be programmed into the QDA4E by selecting the correct pin configuration on the PCB.
- 5. If an address other than that selected for the QDA4E by CONFIGURE command is desired, CONNECT statements must be entered into the SYSCONIF.COM file. SYSCONIF.COM can only be accessed through the system manager's account SYS\$MANAGER. The correct syntax is given in the DEC MicroVMS SYSGEN documentation.

The STARTUP.COM or UVSTART.COM command files in the main system account, SYS\$SYSTEM must not be altered.

6.6.3 Autoconfigure

Autoconfigure is a utility program that finds and identifies I/O devices in the I/O page of system memory. Most devices have a fixed bus address reserved for them. When the computer is bootstrapped autoconfigure polls those addresses—specifically the console status register (CSR) which is usually the first register of the block.

A block of addresses is reserved when a device is detected. The size of the block is determined by the number of registers the device uses. Autoconfigure then looks to the next CSR address space for that same type of device. If

there are no other devices of that type autoconfigure looks to the next valid CSR address. Autoconfigure expects an eight byte block to be reserved for each device not installed in the system. An empty block tells autoconfigure to look to the next valid address space.

Devices with no fixed address are assigned addresses from floating CSR address space. This may be necessary if there are several of the same device in the system. Floating address space is in the vicinity of 76000 to 763776 of the bus I/O page. Devices can also have floating interrupt vector addresses. Floating CSR and interrupt vectors must be assigned in specific sequences depending on the rank of the device (see table below). The presence or absence of floating bus and interrupt vector address devices will affect the assignment of addresses to other floating vector devices.

SYSGEN DEVICE RANKING

Rank	Device	No. of Reg's	Octal Modulus	Rank	Device	No. of Reg's	Octal Modulus
1	DJ11	4	10	17	Res'yd	4	10
2	DH11	8	20	18	RX11 ² 2	4	10
2 3 4	DQ11	4	10	18	RX2112	4	10
4	DU11,	4	10	18	RXV11 ²	4	10
	DUV11						
5	DUP11	4	10	18	RXV21 ²	4	10
6	LK11A	4	10	19	DR11-B3	4	10
7	DMC11	4	10	20	DR11-B ³	4	10
8	DMR11	4	10	21	DMP11	4	10
8 8 8 9	DZ11 ¹	4	10	22	DPV11	4	10
8	DVZ11	4	10	23	ISB11	4	10
8	DZS11	4	10	24	DMV11 ₂	8	20
8	DZ32	4	10	25	DEUNA ²	4	10
9	KMC11	4	10	26	UDA50 ²	2	4
10	LPP11	4	10	27	DMF32	16	40
11	WMV21	4	10	28	KMS11	6	20
12	VMV31	8	20	29	VS100	8	20
13	DWR7Q	4	10	30	TU81	2	4
14	RL11 ² 2	4	10	31	KMV11	8	20
14	RLV11 ²	. 4	10	32	DHV11	8	20
15	LPA11-K	² 8	20	33	DMZ32	16	40
16	KW11-C	4	10	34	CP132	16	40

 $^{^{}m 1}$ DZ11-E and DZ11-F treated as two DZ11s.

- ² The first device of this type has a fixed address while extra devices have floating addresses.
- The first two devices of this type have fixed addresses while extra devices have floating addresses.

An eight byte gap must also be reserved in floating address space for each device type not currently installed in the system. This gap must start on the proper boundary.

A device's CSR address is determined on word boundaries according to the number of bus accessible registers the device has. The relationship of word boundaries and device registers is set out in table below. Autoconfigure only inspects for a device type at one of the possible device boundaries. For instance, autoconfigure will not look for a DMZ32 which has 16 registers at an address that ends in 20.

DEVICE REGISTERS AND WORD BOUNDARIES

Device Registers	Possible Boundaries
1 2 3, 4 5, 6, 7, 8 9 thru 16	Any word XXXXX0, XXXXX4 XXXXX0 XXXXX0 XXXX00, XXXX20, XXXX40, XXXX60 XXXX00, XXXX40

Vector Addresses and Autoconfiguration

Devices are assigned vector addresses in order of rank commencing at 300 (octal) up to 777 (octal). Extra devices of the same type are assigned consecutive vector addresses according to the number of vectors required and starting boundaries for each device type. Table below shows the order of assignment.

The boundaries in the modulus column indicate where vector addresses are assigned. If the modulus is 10 the first vector address for that device must end with a zero (XXO). If the modulus is 4 the first vector must end with with either a zero or four (XXO, XX4).

Vector addresses can only end on an address of four or zero i.e. modulo 4 boundaries (XXO, XX4). If a device has two vectors the first must start on a modulo 10 boundary. Using 350 as a starting point the vectors will be 350 and 354.

FLOATING VECTOR ADDRESS DEVICE PRIORITY RANKING

Rank	Device	Number of Vectors	
1	DC11	2	10
1	TU58	2	10
2 2 2 2 2 3 4	KPTT 1	2 2 2 2 8 2 2 2	10
2	DL11-A1	2	10
2	DL11-B ⁺	2	10
2	DLV11-J ¹	8	40
2	DTATT, DTATT-E.	2	10
3	DP11	2	10
4	DM11-A	2	10
5 6	DN11	1 1	4
6	DM11-BB/BA	1	4
7	DH11 modem control	1 2	4
8	DR11-A, DRV11-B	2	10
9	DR11-C, DRV11	2 4 2 2 2 2 2 2 2	10
10	PA611 (reader + punch	1) 4	20
11	LPD11	2	10
12	DT07	2	10
13	DX11	2	10
14	DL11-C TO DLV11-F	2	10
15	DJ11	2	10
16	DH11	2	10
17	VT40	4	20
17	VSV11	4	10
18	LPS11	6	40
19	DQ11	2	10
20	KW11-W, KWV11	2	10
21	DU11, DUV11	2	10
22	DUP11	2	10
23	DV11 + modem control	3	20
24	LK11-A	2	10
25	DWUN	2	10
26	DMC11	2	10
26	DMR11	6 2 2 2 2 3 2 2 2 2 2 2 2 2 2 2 2 2 2 2	10
27	DZ11/DZS11/DZV11	2	10
27	DZ32	2	10
28	KMC11	2	10
29	LPP11	2	10
30	VMV21	2	10
31	VMV31	2	10
32	VTV01	2	10

33 34	DWR70	2 1 1	10 4
35	RL11/RLV11 ² TS11 ² , TU80 ²	1	4
36	LPA11-K	2	10
37	IP11/IP300 ²	1	4
38	KM11-C	2	10
39	KW11 ₂ C RX11 ²	1	4
39	RX2112	1	4
39	RXV112	1	4
39	RXV212	1	4
40	DR11-W ₂	2 1 1 1 1 1 2 2	4
41	DR11-W2	1	4
42	DMP11	2	10
43	DDV11	2	10
44	DPV11 ML11 ³	2	4
45	ISB11	1 2 2 1 1	10
46	DMV11 ₂	2	10
47		2	4
48	DEUNA ² UDA50 ²	± 1	4
49		8	
	DMF32	8	40
50	KSM11	3	20
51 52	PCL11-B	2	10
52 53	VS100	3 2 1 1 2 2 2 2	4
53	Reversed	1	4
54	KMV11	2	10
55	Reserved	2	10
56	IEX	2	10
57	DHV11	2	10
58	DMZ32	6	20
59	CP132	6	20

System Configuration Example

An example of a system configuration is shown in table below. The configuration includes both fixed and floating addresses and vectors.

KL11 or DL11 have fixed vectors when used as a console.

 $^{^{2}\,}$ The first device has a fixed vector all subsequent device of the same type have a floating vector.

ML11 is a Mass Bus device which connects to the Qbus or Unibus via a bus adaptor.

CSR AND VECTOR ADDRESS EXAMPLE

Controller	Vector	CSR
1 UDA50 1 DZ11 1 UDA50 2 DHV11	154 300 310 320 330	772150 760100 760334 760520

Following table shows the computed CSR addresses and gaps for floating devices. $\,$

FLOATING CSR ADDRESS ASSIGNMENT

		Octal
Installed Device		Address
DJ11	GAP	760010
DH11	GAP	760020
DQ11	GAP	760030
DU11	GAP	760040
DUP11	GAP	760050
LK11A	GAP	760060
DMC11	GAP	760070
> DZ11		760100
	GAP	760110
KMC11	GAP	760120
LPP11	GAP	760130
VMV21	GAP	760140
VMV31	GAP	760150
DWR70	GAP	760170
RL11	GAP	760200
LPA11-K	GAP	760220
KW11-C	GAP	760230
Reserved	GAP	760240
RX11	GAP	760250
DR11-W	GAP	760260
DR11-B	GAP GAP	760270 760300
DMP11		
DPV11	GAP	760310

	ISB11 DMV11		GAP GAP	760032 760330
	DEUNA		GAP	760340
		(ODA 4E)	GAP	7723341
	UDA50	(QDA4E)		
>	UDA50	(QDA4E)		760354
			GAP	760360
	DMF32		GAP	760400
			GAP	760440
	KMS11		GAP	760420
	VS100		GAP	761440
	TU81		GAP	761450
	KMV11		GAP	761460
>	DHV11			761500
>	DHV11			761520
			GAP	761530
	DMZ32		GAP	761540
	CP132		GAP	

¹ indicates a fixed address device

6.7 Cache Operation

6.7.1 QDA4E Disk Cache

The QDA4E implements a disk cache which is designed to facilitate larger and faster data transfers between disk and the host by reducing the time wasted on positioner operations. Even with fast drives 98% of disk time for continuous random access to single sectors of data is taken by positioner operations. The QDA4E offers at least an 80% improvement in access times by reducing the number of disk accesses required. The QDA4E cache is implemented as one megabyte of dynamic RAM.

6.7.2 Read Look-ahead

The user can program the controller to perform read look-ahead in anticipation of impending data requests. The optimum look-ahead value can only be determined within system and application parameters but can range from 0 to 255 blocks. A value of zero will disable the feature. The default value is four.

The anticipated hit ratio for the QDA4E cache is 90% although this can be reduced depending upon the nature of the data accessed. Because most user programs write and read data sequentially there is a high probability that in one fetch operation the controller will be able to satisfy several sequential data reads without the need for further disk accesses.

The cache has been designed to maximize the probability of finding the target data over a range of sequential and non-sequential reference patterns while minimizing cache misses and controller overhead.

6.7.3 Cache Allocation

Cache memory is used to hold the disk cache blocks, a cache map and fixed buffers for special usage. Data from the disk or main memory is stored in blocks at addresses determined by the cache assignment algorithm. Their contents and location are recorded in the cache map.

The cache map consists of a 4-byte entry for each cache block. The cache map is indexed by the cache block number and contains the address (drive number, logical block) of the current occupant together with flags (locked, valid, primary copy not written, shadow copy not written).

Fixed buffers are assigned for RCT buffers (1 per drive) and a single block buffer for disk management I/O. The location and size of all cache variables are held in RAM.

6.7.4 Cache Usage

All disk I/O is done via the cache. A set of cache blocks must be assigned for all transfers and continuous disk operations must be done through continuous cache blocks. Disk and Qbus transfers are performed simultaneously.

6.7.5 Cache Assignment Algorithm

The QDA4E cache implements a contention based hashing algorithm to determine block replacement. A given disk block has a fixed cache address calculated as follows:

- (a) Get the remainder of the logical block number modulo the number of cache blocks.
- (b) Bias this by a fixed offset which is a function of the drive number. (This is so that the same logical block on 2 disks have a different cache block number).

A disk block also has an alternative cache address calculated by biasing it by approximately half the number of cache blocks. The alternate cache block is only used for compare operations.

6.7.6 Cache Operation

The following describes the cache operation algorithm:

Read:

Examine cache for data required.

If all data in cache Transfer data from cache to Qbus

else

Assign cache (lock it and wait if locked already)
Perform read
Unlock cache and flag as valid

Write:

Assign cache (lock it and wait if locked already) Transfer from Qbus to cache and flag as valid Perform write Unlock cache

6.7.7 Cache Disable

The QDA4E cache can be disabled for performance evaluation, engineering, and diagnostic application requirements by selecting that option in WOMBAT. The cache cannot be selectively disabled for a particular drive.

6.7.8 Early Write Notification

The QDA4E implements early write notification where data to be written to disk is retained in the cache and the host is issued a write complete notification. The controller will then write the data to the disk at the most convenient time.

Early Write Notification should be disabled when saving the boot block on the system volume. If it is not disabled, a subsequent reset instruction will clear the cache memory and the boot data will not be written to disk.

It should be noted that in the event of system failure any data residing in the cache will be lost. The early write notification can be disabled by invoking WOMBAT and selecting the appropriate option.

6.8 ESDI Interface

The Enhanced Small Device Interface is an industry standard developed by the ESDI Committee to provide a low cost interface suitable for smaller high performance memory devices. The ESDI interface consists of a 34-pin control cable and a 20-pin data cable.

6.8.1 Control and Data Cables

The control cable allows for a daisy chain connection of up to four drives with the last drive being terminated. The data cable must be attached in radial fashion. The maximum cable length is 3 meters (9.8 feet). Following tables provide control and data cable details.

ESDI CONTROL CABLE (J1) SIGNALS

SIGNAL GROUND PIN PIN SIGNAL NAME -HEAD SELECT 2³
-HEAD SELECT 2 3 5 4 -WRITE GATE 6 7 -CONFIG/-STATUS DATA 8 9 11 -TRANSFER ACK 10 -ATTENTION 12 -HEAD SELECT 2⁰ 14 13 16 -SECTOR/-BYTE CLOCK/-ADDRESS MARK FOUND 15 -HEAD SELECT 2^{1*} 17 18 20 19 -INDEX 21 -READY 22 24 23 -TRANSFER REQ -DRIVE SELECT 1 25 26 27 -DRIVE SELECT 2 28 29 -DRIVE SELECT 3 30 31 32 -READ GATE 33 -COMMAND DATA 34

^{*} Unused but terminated or held false. Ground all odd pins.

ESDI INTERFACE DATA CABLE (J2) SIGNALS

	GNAL GROUN IN PIN	ID
-SECTOR/-BYTE CLOCK/-ADDRESS MARK FOUND -SEEK COMPLETE -ADDRESS MARK ENABLE -RESERVED FOR STEP MODE +WRITE CLOCK -WRITE CLOCK	1 12 3 15 4 16 7 19 8 19	

The four head select lines 2^0-3 allow selection of each individual read/write head with HEAD SELECT 0 being the least significant line. Heads are numbered 0 through 15. Head 0 is selected for removable media drives, Head 0 being selected when all HEAD SELECT lines are high (inactive). The SELECT HEAD GROUP command allows for the addressing of more than 16 heads. Head addressing is continuous from 0. Data can be written to or read from the disk when the WRITE GATE or READ GATE signals are active (low).

6.8.2 Command Data

COMMAND DATA consists of 16 information bits of serial data plus parity. The command data word structure is set out in table below. Flow control is through the handshake signals TRANSFER REQ and TRANSFER ACK with the MSB being transmitted first. The parity bit is odd.

COMMAND DATA WORD STRUCTURE

MOST SIGNIFICANT BIT LEAST SIGNIFICANT BIT

15 14 13 12	11	10	9	8	! ! 7 !	6	5	4	3	2	1	0	P
CMD FUNCTION	CMD	MOD	IFI	ER	! ! !			ALL	ZE	ROS			
CMD FUNCTION !			С	MD	PA	RAM	ETE	R					!

BIT P: PARITY (ODD)

Command Data bits 15 through 12 in combination with the Command Modifier bits 11 through 8 define a variety of ESDI functions. The Command Modifier bits are appended to each of the command functions where applicable. The following table shows the meaning of the various bit combinations. Detailed explanations of these functions can be found in the ESDI Interface Standard documentation.

COMMAND (CMD) DATA DEFINITION

BIT	CT]		12	CMD FUNCTION DEFINITION	CMD MODIFIER APP. BITS 11-8	CMD PARAMETER APPLICABLE BITS 11-0	STATUS CONFIG.DATA RETURNED TO CONTROLLER
0	0	0	0	SEEK	l no	YES	NO
l ŏ	ŏ	ő	ĭ	RECALIBRATE	NO	NO	NO
l ŏ	ŏ	ĭ	ō	REQST STATUS	YES	NO	YES
Ö	ŏ	ī	ĭ	REQST CONFIG	YES	NO	YES
l o	i	0	ō	SELECT HEAD	NO	YES	NO
•	_	_	•	GROUP (OPT)			
0	1	0	1	CONTROL	YES	NO	NO
0	1	1	0	DATA STROBE	YES	NO	NO
1				OFFSET			
0	1	1	1	TRACK OFFSET	YES	NO	NO
1	0	0	0	INITIATE	NO	NO	NO
1				DIAGNOSTICS			
				(OPTIONAL)			
				(OPTIONAL)			
1	0	1	0	RESERVED	-	-	-
1 1	0	1	1	RESERVED	-	-	-
1	1	0	0	RESERVED	-	-	-
1	0	0	1	RESERVED	_	-	-
1	1	1	0	RESERVED	-	-	-
1	1	1	1	RESERVED	-	_	-

6.9 Qbus Interface

All data, address and control information transfers between the processor and disk controller are carried out over the Qbus.

The Qbus consists of 42 bidirectional and 2 unidirectional signal lines wired into the backplane assembly. These are grouped into the following categories:

- * Sixteen multiplexed data/address lines BDAL<15:00>
- * Two multiplexed address/parity lines BDAL<17:16>

- * Four extended address lines BDAL<21:18>
- * Six data transfer control lines BBS7L, BDINL, BDOUTL, BRPLYL, BSYNCL, BWTBTL
- * Six system control lines BHALTL, BREFL, BEVNTL, BINITL, BDCOKL, BPOKL
- * Ten interrupt control and direct memory access control lines BIAKOL, BIAKIL, BIRQ4L, BIRQ5L, BIRQ6L, BIRQ7L, BDMGOL, BDMRL, BSACKL, BDMGIL

Communication is asynchronous, allowing devices with differing data rates to share the bus. A strict master/slave protocol avoids the need for synchronizing clock pulses by implementing handshaking and other control signals between I/O devices.

6.9.1 Interrupts

Interrupt priority for the QDA4E is switch selectable on the PCB. The recommended priority setting is five.

Interrupts suspend program execution while the processor starts the device service routine at a vector address input from the requesting device.

Interrupts are serviced according to device priority. Device priority can be determined in two ways. These are termed 'Position Defined' and 'Distributed' arbitration. Positioned Defined arbitration gives priority to those devices which are electrically closest to the processor. Distributed arbitration implements priority according to the priority levels set on the device hardware. When devices with equal priority generate an interrupt, the processor gives preference to the device which is electrically closest. A previous bus transaction must have been completed before another can be commenced.

The interrupt protocol has three phases:

1. Interrupt Request Phase. The interrupt enable bit in status register is set and interrupt request lines are asserted according to priority settings.

- 2. Interrupt Acknowledge and Priority Arbitration Phase. processor detects the request and checks if any other device with priority is requesting an interrupt. If there are no devices with higher priority seeking an interrupt the processor acknowledges the interrupt.
- 3. Interrupt Vector Transfer Phase. The device outputs vector address bits to the processor which then enters the device service routine.

6.9.2 Direct Memory Access

The QDA4E supports both normal and block mode Direct Memory Access (DMA). During a DMA transfer the processor passes mastership of the bus to the controller.

During block mode DMA transfer the QDA4E has a four microsecond delay after every 16 words to service any pending interrupt or DMA requests from other devices. The QDA4E also detects DMA requests from other devices and will implement a 'DMA Throttle' after eight words. This prevents data loss from other DMA devices which may also share the Qbus.

The QDA4E interleaves address references with bursts of data during DMA. Because the starting memory address is asserted only once every sixteen data words so data throughput is almost doubled.

DMA protocol consists of three phases:

- 1. Bus Mastership Acquisition Phase. The QDA4E requests control of the bus. The processor arbitrates the request then initiates the transfer of bus mastership.
- 2. Data Transfer Phase. The processor provides the controller with the following information utilizing MSCP--block number on the disk, the number of bytes to transfer, and address in main memory, and if the operation is a read or write.
- Bus Mastership Relinquish Phase. Bus mastership is relinquished after completing or aborting the data transfer cycle.

SECTION 7 MAINTENANCE

The QDA4E subsystem is highly reliable, each component providing 30,000 MTBF (Mean Time Between Failure) and a typical MTTR (Mean Time To Repair) of a half-hour.

All drives are formatted by System Industries and should perform on startup.

Preventive maintenance consists of keeping components clean and free of dust.

7.1 Diagnostics and Troubleshooting

If you wish to run a quick verification of system, perform read and write transfers to each disk as well as creating and deleting a directory. If you can read but not write, check for write protection status.

Check cabling for attachment and connector looseness.

Verify that system has been configured for the proper number of drives and check drive and controller jumpers and switches.

Run WOMBAT as described in the QDA4E controller section.

7.2 Field Replaceable Units

There are no FRU's except for the component itself: the QDA4E, SI55, SI56, SI57, DA123, 5 1/4 Inch Quick Release Tray and Package, or cables.

7.3 Removal and Replacement Precautions

Follow the rules listed below to prevent possible injury to yourself or the equipment:

o Power down disk drive--subsystem--and disconnect all AC power cords to the cabinet, disk drive, or DC supply before performing any work.

WARNING

NEVER REMOVE OR INSTALL ANY PCB OR DISCONNECT ANY CONNECTOR, PLUG, OR WIRE WHILE POWER IS ON.

- o When removing connectors, do not grasp the wire bundle which can damage the cable. Rather, hold the connector firmly by its side and pull out.
- o Do not remove any parts that are not specified in the replacement procedure.
- o Package printed circuit boards in special electrostatic free envelopes.
- o When working near printed circuit boards, ground yourself with an antistatic wrist strap.